Probabilistic Causation, Preemption and Counterfactuals

PAUL NOORDHOF

Counterfactual theories of Causation have had problems with cases of probabilistic causation and preemption. I put forward a counterfactual theory that seems to deal with these problematic cases and also has the virtue of providing an account of the alleged asymmetry between hasteners and delayers: the former usually being counted as causes, the latter not. I go on to consider a new type of problem case that has not received so much attention in the literature, those I dub *catalysts* and *anti-catalysts*, and show how my account needs to be adjusted to deliver the right verdicts in these cases. The net result is a particular conception of a cause that I try to spell out in the closing section of the paper. In that section, I also briefly discuss causal asymmetry and the purpose behind providing a counterfactual theory of causation.

In this paper, I put forward a counterfactual theory of causation. I develop it by discussing some difficulties, which have plagued previous accounts of this type:

(I) Probabilistic Causation
(II) Early and Late Preemption
(III) Hasteners and Delayers.

I will begin by outlining the main problem that a counterfactual theory faces—its treatment of preemption, especially probabilistic preemption—and discussing some earlier, inadequate, attempts to deal with it, in particular the benchmark theory provided by David Lewis. I will use this discussion as a springboard to motivate my own theory. To some extent, my theory will be justified by its successful treatment of various problem cases. However, getting the right results is not really enough, and so I try to explain how the problem cases highlight various features of our notion of causality and how the various clauses of my account capture these features. To close, I will discuss the extent to which my theory can be considered to illuminate of the nature of causation.

I am going to be concerned with causal relations between particulars—specifically events. I have tried to keep my commitments at a minimum with regard to their individuation in the hope that the theory put forward could be coupled with any account of their nature. I also hope that, although I have selected an event ontology for presenting my theory, it could be adapted to any other preferred account of the relata of causation. In order to keep an already large paper manageable, I'm afraid these will have to remain hopes.

Mind, Vol. 108 . 429 . January 1999

1. Probabilistic causation and preemption: a devastating combination

Modern physics has convinced many that probabilistic causation is possible. One illustration of this is radioactive decay (taken from Mellor 1995, pp. 52–3). Radium and uranium have radioactive isotopes which turn into other elements—*decay*—when subatomic particles such as α or β particles "tunnel" out of the nuclei of these elements. The decay is not governed by deterministic laws. At time *t* (say), it is not settled that an atom of radium starts to decay. Rather the laws which govern decay give the atom a particular less than 1 chance of decaying during a time interval. However, if the nucleus of an atom is bombarded by a subatomic particle, then it is almost certain to decay. Consequently, it seems right to say that the bombardment of the nucleus of the atom is a cause of the decay. Yet, since the chance of decay is not 1, the bombardment is not sufficient for the decay. And since the decay might have happened anyway without the bombardment, the bombardment is not necessary—that is, it is not even necessary in the circumstances we are envisaging.[1]

If probabilistic causation is possible, counterfactual theories cannot appeal to the following subjunctive conditionals:

(i) If e_1 were not to occur, then e_2 would not occur

(ii) If e_1 were to occur, then e_2 would occur

in order to capture the notion of counterfactual dependence between e_1 and e_2 in terms of which the causal relation is then defined. The *joint* holding of these conditionals imply that e_1 is both necessary and sufficient for e_2 to occur. The existence of probabilistic causation is incompatible with this even if the ancestral of the relation of counterfactual dependence—rather than the relation of counterfactual dependence itself—is taken to be the causal relation (see Lewis 1973, p. 167).[2] Such an account would still require that, if there were no intermediary between e_1 and e_2, (i) and (ii) would hold. For a case of probabilistic causation, they would not. Nor does taking the ancestral of counterfactual dependence introduce indeterminism for cases of mediate causation.

To deal with this, David Lewis put forward an account of probabilistic dependence which may be formulated as follows.

[1] It is this fact that makes Ramachandran's M-Set Analysis of Causation unable to handle probabilistic causation—contrary to advertisement (Ramachandran 1997, pp. 272–3, 276—see Noordhof 1998a, pp. 459–60).

[2] Some doubts have been expressed over whether such counterfactuals can capture this necessity and sufficiency (e.g. Mellor 1995, pp. 28–30). However, there is reason to suppose that these doubts can be assuaged (Noordhof 1998b).

> Event e_2 *probabilistically-depends* on a distinct event e_1 iff it is true that: if e_1 were to occur, the chance of e_2's occurring would be at least x, and if e_1 were not to occur, the chance of e_2's occurring would be at most y, where x is much greater than y. (Lewis 1986, pp. 176–7)

The phrases "at least" and "at most" have been introduced to try to accommodate Lewis's point that in the closest e_1-worlds—and also in the closest not-e_1-worlds—the chance of e_2 may fluctuate so that there is no precise chance that e_2 has. The chances mentioned are objective, single case chances as opposed to frequencies, features of the world and not credences (Lewis 1986, pp. 177–8). He remarks that "much greater than" is to be understood here as "by a large factor" not "a large difference" (Lewis 1986, pp. 177–8). This is to accommodate the fact that both probabilities may be small. He then defined causation by taking the ancestral of probabilistic dependence. As a result, we have

> For any actual distinct events e_1 and e_2, e_1 causes e_2 iff there are events $x_1, \ldots, x_n$ such that x_1 probabilistically depends upon $e_1, \ldots$, e_2 probabilistically depends on x_n. (Lewis 1986, p. 179)[3]

For the discussion that follows, the most important feature of Lewis's theory is his subsequent remarks about how we should understand the chance that a particular event has of occurring. The chance of an event varies with time, having "different chances at different times before it occurs" (Lewis 1986, pp. 176–7). In assessing whether an event e_1 is a cause of e_2, Lewis suggests that we take the chance of e_2—p(e_2)—to be that which it had *immediately after e_1 occurs or fails to occur* (Lewis 1986, p. 177). Counterfactuals with probabilities in their consequents are, then, to be assessed in roughly the following manner. If the antecedent is false, hold fixed the history of the world as much as possible up to just before the circumstances mentioned in the antecedent fail to hold—in the present case this will be either that e_1 occurs or that it fails to occur—and change the world just enough to make the antecedent true. In these circumstances, assess the chance of e_2 just after e_1 occurs (or fails to occur).

Lewis's account would capture the intuition that the bombardment by the subatomic particle is a cause of the decay. The bombardment made the chance of an atom of the element decaying very much greater than the chance would have been had no bombardment taken place. However,

[3] I'm suppressing for the moment Lewis's reformulation to deal with late preemption. Lewis himself took the present account to deal with early preemption, of which Figure 1 is an example, and the reformulated account does not in fact help with the difficulty identified by Menzies (see Lewis 1986, pp. 200–2; Menzies 1989, pp. 645–6).

Peter Menzies offered a counterexample to Lewis's proposal (Menzies 1989, pp. 645–6). Consider the diagram below.

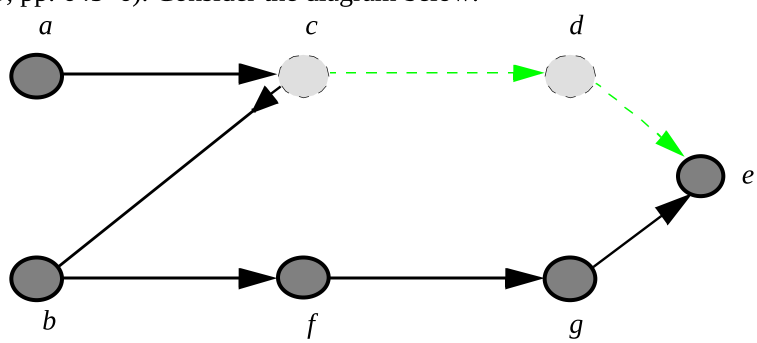

Figure 1

The lettered circles are neurones which either fire or don't fire (the broken circles). The connections stimulate (forwards arrow) a neurone to fire or inhibit (backwards arrow) a neurone from firing. A neurone that is both stimulated and inhibited does not fire. The process leading from *b* to *e* is unreliable. The inhibitory axon between *b* and *c* and the process leading from *a* to *e* are more reliable—the *a*–*e* process very much more reliable left to itself. The problem is that Lewis's account would imply that *e*'s firing probabilistically depends upon *a* firing in spite of the fact that *c* and *d* don't fire. Just after *a* fires, it is possible that *b*'s firing won't inhibit *C* from firing and hence that the much more reliable *a*-chain run to completion. So, at the point in time just after *a* fires, the chance of *e* firing later is very much greater than the chance would have been at that point in time if *a* had not fired (Menzies 1989, pp. 647, 653). The conditions Lewis identifies are not sufficient for causation.

What else is needed? I want to discuss two proposals that won't work as they stand but which suggest an account which I believe will. Both attempt to characterise the way in which the preempted process has not run to completion whereas the preempting process has. Peter Menzies suggests that

> e_1 causes e_2 only if there is a chain of unbroken causal processes running from e_1 to e_2. (Menzies 1989, p. 656, the formulation has been adjusted to fit my discussion)

The idea is that for any finite sequence of times $<t_1, \ldots, t_n>$ between the time of e_1 and e_2, there is a sequence of actual events occurring at these times $<x_1, \ldots, x_n>$ where x_1 is probabilistically dependent upon e_1, … e_2 is probabilistically dependent on x_n. Call this an unbroken causal process. A finite sequence of events <*a*, *b*, *c*, … > is a chain of unbroken causal processes if and only if there is an unbroken causal process running from *a* to *b*, an unbroken causal process running from *b* to *c* and so on. Talk of *chains* of unbroken causal processes is necessary to deal with the fact that e_1 can be a cause of e_2 even if there are some sequences of events between e_1 and e_2 which may not pairwise probabilistically depend upon each other. One example would be the finite sequence of events which just includes *b*'s firing and *e*'s firing in the original diagram (see Menzies 1989, pp. 654–5; Menzies 1996, pp. 93–4). Menzies's theory allows *b*'s firing to be a cause since unbroken causal processes can be patched together between *b*'s firing and *e*'s firing whereas no chain of unbroken causal processes can be patched together between *a*'s firing and *e*'s firing.

Unfortunately, Menzies's account is inadequate—as he now recognises. First, it rules out temporal action at a distance. It insists that there must be events at all the times between e_1 and e_2 for e_1 to cause e_2. Any theory which failed to rule this out a priori would have an advantage (Menzies 1996, p. 94). Second, it cannot handle cases of either deterministic or probabilistic late preemption, that is cases in which the process preempted is preempted by the occurrence of the effect.[4] Consider the diagram below.

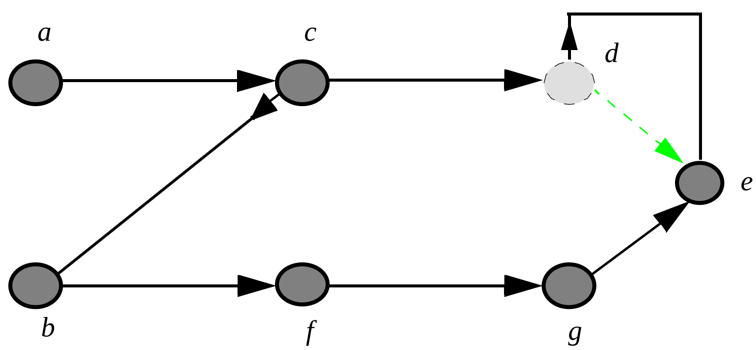

Figure 2

As before, the *a*–*e* process is very reliable whereas the *b*–*e* process is unreliable. The crucial difference is that it is *e*'s firing which inhibits *d* from firing. If *e*'s firing had not occurred at the time it did as a result of the *b*–*e* process it would have occurred later—and hence after *d* firing—as a result of the *a*-chain.[5] The problem is that it is hard to see how Menzies's proposal could obtain the result that *a*'s firing is not a cause. Since *d*'s fir-

[4] Menzies has acknowledged that this is so for the deterministic case (Menzies 1996, pp. 95–6). I focus on the probabilistic case which he does not discuss. It raises an important issue for my subsequent discussion. Thanks to an anonymous referee for setting me straight on this.

[5] The case obviously rests upon the assumption that particular events don't have their times of occurrence essentially but could occur later than they did. This matter is discussed later. For a defence of the view, see Lewis (1986, pp. 204–5, 249–50).

ing occurs after *e*'s firing, at all times up until the time of occurrence of *e*'s firing, there will be events in the *a*-chain upon which *e*'s firing probabilistically depends. It is only if we consider times after *e*'s firing occurred—but before it would have occurred if *e*'s firing had not been brought about earlier by the *b*-process—that we find a missing event: *d*'s firing.

What do these problems show about Menzies's account? I think he was right to suppose that the difference between pre-empter and pre-empted should be captured in terms of whether the causal chain between putative cause and effect is complete or not. It is just that he had the wrong account of what makes a process complete. We don't need a chain of unbroken causal processes involving events at every moment in time between putative cause and effect. We just need all the events which were actually necessary for the causal chain to be present. Understanding completeness Menzies's way resulted in ruling out action at a distance and the problem with the late preemption case. But his way is not mandatory as I shall try to demonstrate.

A second idea that might capture the notion of a causal process being complete concerns the time at which the probability of the effect is assessed. The thought is that Lewis identified the wrong time at which p(*e*) ought to be assessed. He suggested that it ought to be just after the candidate cause, *a*'s firing, occurred. But it should be assessed just before the effect occurred. In the case of early preemption (Figure 1), just before *e*'s firing occurred, elements of the *a*–*e* process would have failed to occur, namely the firing of *c* and *d*, so the occurrence or non-occurrence of *a*'s firing would not be relevant to p(*e* fires). This would reveal the fact that the *a*-chain is incomplete. Unfortunately, once again, this delivers the wrong result in cases of late preemption (Figure 2). The problem is that *e*'s firing occurs prior to the preempted event in the *a*-chain. So at the time just before *e*'s firing, everything is in place for *a*'s firing to raise the probability of *e* firing. It is only after the firing of *e* occurs that things begin to go wrong. So changing the time at which we assess the chance of an effect *by itself* won't deliver the right result in cases of late preemption.

2. Late preemption, quasi-dependence and an alternative analysis of completeness

My account of what makes a process complete is best understood by considering a closely related approach upon which it tries to improve. To deal with cases of late preemption, Lewis introduced a notion of quasi-dependence, which may be formulated as follows.

Event e_2 quasi-depends on a distinct event e_1 iff

either (a) there are events $x_1, \dots, x_n$ such that x_1 counterfactually depends upon $e_1, \dots, e_2$ counterfactually depends on x_n,

or (b) the intrinsic character of the process involving e_1 and e_2 is just like that of processes in other regions in this or other worlds with the same laws and in the great majority of these regions—measured by variety—these processes satisfy (a). (Lewis 1986, p. 206)

This proposal does manage to deal with most cases of deterministic late preemption. Consider a deterministic version of the setup in Figure 2. *b*'s firing would not satisfy clause (a) of this account because there is no event between *b* firing and *e* firing upon which *e*'s firing probabilistically depends. Take the firing of *g* as an illustration. If *g* were not to fire, *e* would still fire brought about by the firing of *d*. However, *b*'s firing would be a cause of *e*'s firing if it satisfied clause (b) and there is every reason to think that it will. In the great majority of circumstances, chains with the same intrinsic character as the *b*-chain will occur without chains of the same type as the *a*-chain. In these circumstances, Lewis claims, there will be a chain of counterfactual dependencies of the kind to satisfy clause (a).[6]

There are two significant problems for this approach. One occurs in the case of indeterminism where we substitute probabilistic dependence for counterfactual dependence in clause (a). As we have already noted, *e*'s firing probabilistically depends upon *a*'s firing whether or not the *b*-chain is present. So *a*'s firing satisfies the first clause of Lewis's definition of quasi-dependence and hence the firing of *a* is proclaimed a cause. A second reason for feeling unhappy about Lewis's approach is that it rules out brute singular causation in cases of preemption (Ganeri, Noordhof and Ramachandran 1996, pp. 223–4). We have a case of brute singular causation if e_1 causes e_2 without any backing regularity. If this is allowed generally, then there is no reason to disallow it in cases of late preemption. Suppose the *b*-chain is a case of brute singular causation. Then in the great majority of circumstances, intrinsically similar—i.e. *b*-type—processes may not stand in a chain of counterfactual dependence at all. The *b*-chain was a one-off. So the firing of *b* wouldn't be a cause by Lewis's account. Of course, if there were no other way to capture the intuitions we have about what are causes and what are not in cases of late preemption, then we might have to adopt this strategy. However, there is an alternative.

[6] Lewis dismisses cases in which the two processes nomically co-occur as "very peculiar indeed" and suggests that they are spoils to the victor (Lewis 1986, p. 207).

Hence, it is a legitimate criticism of Lewis's approach that it appears gratuitously to rule out brute singular causation—something that Lewis himself has said he is loath to do (Lewis 1986, p. 169).

The alternative is the one I put forward along with Jonardon Ganeri and Murali Ramachandran (Ganeri, Noordhof and Ramachandran 1996; Ganeri, Noordhof and Ramachandran 1998). It was suggested by focusing on what Lewis might have had in mind by appealing to the notion of intrinsic similarity of processes in distinguishing between the pre-empting and pre-empted processes. Suppose that both the *b*-type-chains and the *a*-type-chains occur independently of the other in the great majority of circumstances. When they do, both processes involve chains of counterfactual dependence up to and including an event of the same type as the firing of *e*. What led Lewis to suppose that only the *b*-type processes were intrinsically similar? In this case, the answer seemed to be that when *a*-type processes occurred independently of *b*-type processes, the *a*-type processes had an event that they did not actually have in the late preemption case: *d*'s firing (Ganeri, Noordhof and Ramachandran (1996, p. 220).

Having identified the work the appeal to intrinsic similarity was doing in this case, we wondered whether it was possible to formulate a counterfactual account of deterministic causation which dealt with the preemption cases but which allowed for one-off brute singular causation. Instead of appealing to the existence or otherwise of intrinsically similar processes to ground the asymmetry between preempting and preempted processes, we considered whether there would be an asymmetry if we considered these particular preempting and preempted processes in appropriate counterfactual circumstances (an idea present in Honderich 1988, p. 19). The account at which we eventually arrived was this.

> For any actual, distinct events e_1 and e_2, e_1 causes e_2 iff there is a (possibly empty) set of possible events Σ such that
>
> (I) e_2 is Σ-dependent on e_1, and,
>
> (II) every event upon which e_2 Σ-depends is an actual event.[7]

Where

> For any events e_1 and e_2, and any set of events Σ, e_2 Σ-*depends* on e_1 iff

[7] This is a slightly modified version of the account put forward in Ganeri, Noordhof and Ramachandran (1998) in response to Byrne and Hall's criticisms (see Byrne and Hall 1998). It adjusts the account to accord with my conviction that causation is not transitive due to the counterexamples offered by McDermott—mentioned in passing in the article from which this formulation is taken (see McDermott 1995, pp. 531–3). I have also adjusted the lettering to fit with my discussion.

(i) if neither e_1 nor any of the events in Σ were to occur, then e_2 would not occur,

and

(ii) if e_1 were to occur without any of the events in Σ, then e_2 would occur.[8]

It might look (over) complicated but the idea is simple enough. Roughly, the Σ-set mechanism was a means by which we could formalize the introduction of appropriate counterfactual circumstances, the "actual event" clause—clause (II)—was the asymmetry to which we appealed to distinguish preempted from preempting causal chain. Let me run through a few examples to underline the point. Suppose I shoot a man dead before the poison you gave him took effect. I try to claim that I didn't really kill him because if I hadn't shot him, he would have died anyway (from your poison). Your natural response would be

> But if you had shot him without my having poisoned him (putting this event in Σ), he would still have died (satisfying clause (ii) of Σ-dependence). So you did enough to kill him by yourself. If neither you had shot him nor I had poisoned him, then he would still be living (satisfying clause (i) of Σ-dependence). I admit the same could be said of my poisoning. That makes us both potential causes of his death (satisfying clause (I) of the definition of causation). What makes you the cause and me not is that the causal chain from the administration of the poison to death is *incomplete*. There are missing events—those which would have occurred when the poison took effect—that did not occur because the man was shot dead in the meantime. If the poison had been administered without your shooting him, then there would have been events in the chain up to the man's death which did not occur in the actual circumstances (so failing clause (II) of the definition of causation). The same cannot be said of your shooting (so it passes clause (II) of the definition of causation). That's why you killed the man and not me.

As another illustration, consider the example given by Figure 2—taking it to show a deterministic case of late preemption rather than the indeterministic case I described earlier. If we put *a*'s firing in Σ, then *e*'s firing becomes Σ-dependent on *b*'s firing. Similarly, if we put *b*'s firing in Σ, then *e*'s firing becomes Σ-dependent on *a*'s firing. So both are candidate

[8] In the original piece, we appealed to a might-conditional to characterise the second condition. This was to deal with the following kind of case. Suppose Patel is not a marksman. Although he fires and hits a balloon, it was just a fluke. That suggests in some of the closest worlds in which he fires, he misses. However, Patel's actual shot still caused the balloon to burst. We emphasised that that's because the bursting of the balloon *might* occur as a descendent of the shooting (see Ganeri, Noordhof and Ramachandran 1996, pp. 221–2). I am suppressing this point because it is unnecessary to introduce might-conditionals if one starts appealing—as I am about to do—to probabilities.

causes. The difference is that in the first case—in which we are assessing whether *b*'s firing is a cause—*e*'s firing is only Σ-dependent on actual events. It is not Σ-dependent on the firing of *d*, for instance, because if neither *d* nor *a* had fired (*a*'s firing being the event in Σ), *e* would still have fired as a result of *b*'s firing. By contrast, when we turn to assess whether *a*'s firing is a cause, we find that with *b*'s firing in Σ, *e*'s firing is Σ-dependent on *d*'s firing. But *d*'s firing is not an actual event. Hence *a*'s firing is not a cause (it fails clause (II)).

This account seems to work well for deterministic cases. It has the advantage of failing to rule out brute singular causation. It also doesn't rule out temporal action at a distance since it doesn't require that there is an event at every moment in time on the chain between e_1 and e_2. All it relies upon is the existence of possible events which were actually suppressed due to preemption. Unfortunately, it is not so good when we turn to cases of indeterministic late preemption. The natural way to extend it is to substitute a notion of probabilistic Σ-dependence, defined as follows.

> For any events e_1 and e_2, and any set of events Σ, e_2 *probabilistically* Σ-*depends* on e_1 iff
>
> (i) if e_1 were to occur without any of the events in Σ, then p(e_2) would be at least *x*
>
> (ii) if neither e_1 nor any of the events in Σ were to occur, then p(e_2) would be at most *y*
>
> (iii) $x >> y$.

But weakening Σ-dependence to probabilistic Σ-dependence causes problems.

Consider first the case represented by Figure 1. The suggestion was that *b*'s firing satisfies the definition of causation if *a*'s firing is put in Σ. But in the case of indeterminism, it is not so clear. The problem is that even if *a*'s firing did not occur, it is still possible that *c*'s and *d's* firing might occur. Since these are still options just after *b*'s firing, there is no guarantee that *b*'s firing raises the probability of *e*'s firing sufficiently to count as a cause. However, this is only a problem because we have not been sufficiently clear about the time at which p(*e*'s firing) should be assessed. The problem arises if we take the time of assessment to be just after the firing of *b*. So instead, let us take the time of assessment of p(*e*'s firing) to be just before *e* fires, the option we considered at the end of the last section. By then, it will be clear that *c* and *d* have not fired. So the background probability of *e* firing will be low. The firing of *b* will be a cause by the definition because, against this background, it very much raises the chance of *e* firing.

Consider now the application of the "actual events" clause. In Figure 1, the firing of *e* will be probabilistically Σ-dependent on the firing of *d* whatever one puts into (or leaves out of) the Σ-set.[9] This would seem to rule out the firing of *b* being a cause since it would fail the "actual events" clause due to the probabilistic Σ-dependence on firing of *d* even though, intuitively, this event has nothing to do with the *b*'s firing causing *e*'s firing—not a result we want. However, in fact there is no problem. All we have to do is put the firing of *d* as part of the Σ-set for demonstrating that the firing of *b* satisfies my definition. In that case, the firing of *e* would not probabilistically Σ-depend upon the firing of *d*. The conditionals we would have to consider would be these.

> If *d*'s firing were to occur without the firing of *d* (or the rest of the Σ-set) occurring, it would be the case that the chance of *e* firing would be at least *x*.
>
> If neither *d*'s firing nor *d*'s firing (nor the rest of the Σ-set) were to occur, it would be the case that the chance of *e* firing would be at most *y*.

Since the antecedent of the first conditional is necessarily false, we could put any value for *x* including 0 and the conditional would be true. So the conditionals fail to proclaim (except trivially) that *e*'s firing is probabilistically Σ-dependent on *d* firing. We can rule out trivial satisfaction. So the firing of *b* does pass clause (II). This strategy will apply quite generally. By contrast, *a*'s firing will not be a cause of *e*'s firing because if one puts one of the non-actual events in the chain (e.g. the firing of *d*) into Σ, *e*'s firing would not probabilistically Σ-depend on the firing of *a*. So even though the firing of *a* would pass clause (II) with this Σ-set, it would not pass clause (I).

A more substantial difficulty with this characterisation of probabilistic Σ-dependence is revealed by the following case.

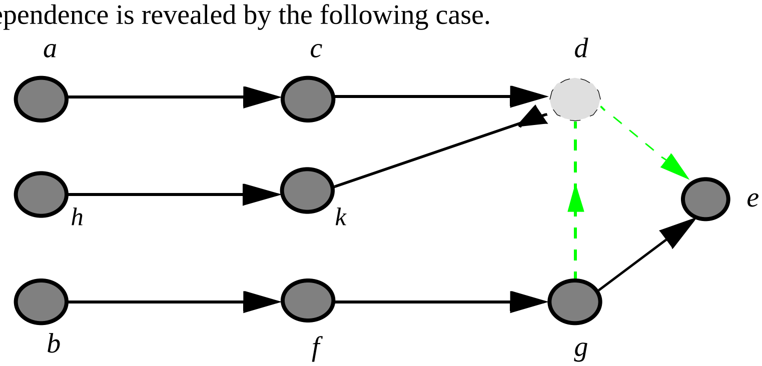

Figure 3

[9] This is not so in a deterministic version of this case since if the firing of *g* is not put in Σ (which it won't be if we are testing to see whether *b*'s firing is a cause), *d*'s firing can't raise the chance of *e*'s firing. P(*e* fires) would already be at 1.

Here, the *a–e* process is very reliable overall but it has a weak link: the *d–e* connection. This connection is much less reliable than the corresponding connection on the *b–e* chain: *g–e*. The *d–g* inhibitory axon is very reliable. The *a*-chain also has one particularly strong link: the *c–d* connection. If *k* had not fired to inhibit *d*, then the chance of *d* firing given that *c* fired is 1. The *b–e* chain is very unreliable. The only strong link is the *g–e* connection. What happens is that *a*, *h* and *b* fire; *d* fails to fire because of the inhibitory *h*-chain; so *e* fires because of the *b*-chain.

Intuitively, we want to say that *a*'s firing is not a cause of *e*'s firing. However, *a*'s firing would satisfy the "probabilistic Σ-dependence" clause with *h*'s firing in Σ. If *h* did not fire, then *d* fired. In these worlds, *a*'s firing would have raised the chance of *e* firing just before *e*'s firing occurred. This makes *a*'s firing a potential cause. The question is whether *e*'s firing is probabilistically Σ-dependent on *d*'s firing so that the "actual events" clause rules *a*'s firing out from being a cause. It seems that the answer is no. If *d* were to fire, then it is still rather unlikely that *e* would fire as a result of the d-e connection—the *d–e* connection being so unreliable—and *d*'s firing is very likely to inhibit *g* from firing. This means that the firing of *d* would inhibit the far more sure way of bringing about the firing of *e* at this point in time, namely by the *g–e* connection. On the other hand, if *d* were not to fire, *g* would not be inhibited from firing and *e* would be likely to fire as a result. So overall the firing of *d* would lower the chance of *e* firing. The firing of *e* does not probabilistically Σ-depend upon *d* firing. As a result, the "actual events" clause would not exclude the firing of *a* from being a cause.

A natural diagnosis of what's gone wrong is that, in searching for a Σ-set to reveal whether an effect is probabilistically Σ-dependent on a particular candidate cause, we are not concerned with whether the intermediate links in the chain are also revealed to be candidate causes. Obviously, if we have a causal chain, the effect can be shown to probabilistically Σ-depend upon the intermediate links too. But there is no reason to expect that a Σ set for the candidate cause will also be sufficient for the intermediate links. So having arrived at a Σ-set which makes e_2 Σ-dependent on e_1, we should consider arbitrary additions to the Σ-set which do not undermine the connection between e_1 and e_2 by including events in the chain connecting them but which may reveal that e_2 is Σ-dependent on other events in the chain. If there is a genuine causal chain, we will be able to add to this Σ-set any number of events (to yield a superset, Σ*) so as to reveal how the effect is probabilistically Σ-dependent on each of the events yet none of these events is non-actual. The following clause captures the idea.

(II)′ For any superset of Σ, Σ*, (where $\Sigma \subseteq \Sigma^*$), if e_2 probabilistically Σ*-depends upon e_1, then every event upon which e_2 probabilistically Σ*-depends is an actual event.[10]

For instance, in the case described in Figure 3, all we would have to do is consider what would happen if we added *g*'s firing to the events in Σ to get Σ*. The firing of *e* would probabilistically Σ*-depend upon *a*'s firing. So the connection between *a*'s firing and *e*'s firing would not have been disturbed. But now *e*'s firing would probabilistically Σ*-depend upon *d*'s firing. The firing of *d* would no longer lower the chance of *e*'s firing by inhibiting *g* from firing since it is already settled that *g*'s firing is not to take place. So the firing of *a* is ruled out as a cause.

This has rectified the consequences of introducing *probabilistic* Σ-dependence for the formulation of clause (II). Unfortunately, a substantial problem remains for this formulation of probabilistic Σ-dependence. It arises regarding our old friend the probabilistic late preemption case—Figure 2. *a*'s firing can satisfy clause (I) of the account "the probabilistic Σ-dependence clause" with no events in Σ. The question is whether the "actual events" clause can indicate what is wrong with counting *a*'s firing as a cause. There is reason to think not. What we need is for *e*'s firing to probabilistically Σ-depend upon *d*'s firing. Then, since *d*'s firing is a non-actual event, *a*'s firing would fail to be a cause. While it is true that the chance of *e* firing is pretty high if *d* has fired, its chance also seems pretty high if *d* were not to fire. To see this consider first what is the case if we assess p(*e* firing) in the way that Lewis recommends—i.e. just after *d* fires or fails to fire. Then, since *d*'s non-occurrence is settled after *e* has fired (inhibiting *d* from firing), we will be assessing the probability of a past event. Generally, these are assumed to have a probability of 1. However, even if this assumption is challenged, p(*e* firing) will be pretty high. Given that the *a*–*d* chain is very reliable, if *d* has not fired, this is more likely to be due to its firing being inhibited by *e* firing than it is due to the simple failure of the *c*–*d* connection. Things are no better if we assess p(*e* firing) in the way I have recommended, namely just before *e* fires. The same reasoning seems to apply. We would have to assess p(*e* firing) just before *e* fires given that it is settled that *d* does not fire. That would once more suggest that *e* is likely to have fired. If the reasoning does not apply, then that would surely be because just before *e* firing, the firing or otherwise of *d* is not settled. In this case, there will be no change in p(*e* fires) at that point due to *d* firing or failing to fire in the future. Either way, *e*'s firing is not probabilistically Σ-dependent upon *d* firing.

[10] Although Murali Ramachandran and I developed this proposal together, he deserves the credit for its final formulation reproduced here.

There does not seem to be any Σ*-set which definitely changes this verdict. Let us make Σ* = {*b*'s firing, *f*'s firing, *g*'s firing}. The question is whether this will be enough to make *e*'s firing probabilistically Σ*-dependent upon *d*'s firing. It is hard to see why. The firing of *e* may have a background probability greater than 0 of occurring when it did even with the absence of the *b*-chain. We noted this in our original case of indeterminism. And, if *e* does have some background probability of firing and, in the circumstances we are considering did fire, then Lewis's method of assessing counterfactuals by keeping the past history up to the circumstances envisaged in the antecedent as similar as possible would suggest that we retain the firing of *e* at the time it did. But then the same reason I used to explain why there is no probabilistic Σ-dependence applies here too. The failure of *d* to fire increases the likelihood that *e*'s firing inhibited *d*'s firing. So an application of the "actual events" clause would not rule out the firing of *a* from being a cause.

3. Hasteners, delayers and probabilistic dependence

We need a change of tack. Consider once more our troublesome case of late preemption. One reason for thinking that *a*'s firing is not a cause is that *d*'s firing does not occur. But another reason for thinking that *a*'s firing is not a cause is that, even though the firing of *a* raised the probability of *e* firing, it did not raise the chance of *e* firing *at the time e fired*—only later (as a result of *d*'s firing). This might suggest the following account of probabilistic Σ-dependence which I shall call "probabilistic Σ-time-dependence".

e_2 probabilistically Σ-time-depends upon e_1 if and only if

(1) If e_1 were to occur without any of the events in Σ, then it would be the case that $p(e_2 \text{ at } t) \geq x$

(2) If neither e_1 nor any of the events in Σ were to occur, then it would be the case that $p(e_2 \text{ at } t) \leq y$

(3) $x >> y$

(4) e_2 occurs at t.

The proposal appears to get the right answer in the case of late preemption. *a*'s firing fails to be a cause because, although *a*'s firing may raise the chance of *e* firing at some time (later than it did), it does not raise the chance of *e* firing at the time it fired (as a result of *b* firing). By contrast, *b*'s firing is a cause of *e* firing because it did raise the chance of *e* firing when it did.

Unfortunately, this proposal is no good. There are cases of preemption I will discuss in the next section for which it delivers the wrong answer. However, there are some more straightforward cases which nicely highlight another difficulty. Suppose that a forest fire occurs in June as a result of electrical storms. These electrical storms occurred for the previous two months. But there was heavy rain in April and the vegetation had been too damp to catch fire until June. If the heavy rain had not occurred in April, the forest would have caught fire in May. It should be clear that, if we just appealed to the analysis of probabilistic Σ-time-dependence offered above in an analysis of causation, we would have to conclude that heavy rain in April was a cause of the forest fire and not just the forest fire occurring in June. It raised the chance of a forest fire in June by delaying it from May. If the fire had occurred in May, the forest would have burnt down and there would have been nothing to catch fire in June (the example comes from Bennett 1987, pp. 373–4). Equally, suppose that a doctor's treatment delayed the onset of a particularly unpleasant phase of an illness but, because the treatment was not entirely effective, it occurred after all at the end of the second week of illness. Again an analysis of causation just based on the notion of probabilistic Σ-time-dependence would make the doctor's treatment a cause of the unpleasant phase of the illness and not just a cause of its occurrence at the end of the second week. I think that in both these cases, this would be the wrong diagnosis.

The usual assessment of why the April rains and the doctor's treatment aren't causes is because they are delayers of their alleged effects. However, this does not go deep enough. It is not clear how, by merely delaying an event, another event should be disqualified from causing it. As we have seen, what matters in general is whether the putative cause raises the chance of another event's occurrence given certain conditions. It is here that we should look for the basis of our intuitions. A preliminary diagnosis of what is wrong with both the April rains and the doctor's treatment being causes in the situations envisaged is that, while they raise the chance of their alleged effects occurring at a particular time, they don't make their effects more probable per se.[11] That suggests the following generalisation of the idea of probabilistic Σ-time-dependence.

e_2 probabilistically Σ-depends upon e_1 if and only if

[11] This is on the assumption that the time events actually occur is not essential to their identity. If it were essential, then an analysis appealing to probabilistic Σ-time-dependence would work. However, even if we fix matters this way at the level of metaphysics, it does not capture the way we talk. We say that particular events and not just particular types of events might have occurred earlier or later than they did. We would still need to characterise what should hold between these artifacts of our thought and talk. The rest of my theory might be taken to address that task by the ardent exponent of time of occurrence being essential to event identity.

(1) If e_1 were to occur without any of the events in Σ, then for some time t, it would be the case that, just before t, p(e_2 at t) $\geq x$

(2) If neither e_1 nor any of the events in Σ were to occur, then for *any* time t, it would be case that, just before t, p(e_2 at t) $\leq y$

(3) $x >> y$.[12]

The intuitive idea is that, relative to the events in Σ, the presence of a cause does not just make the probability of an event at a time very much greater than it would be if the cause were not present. It makes the probability of the event at that time greater than the probability of that event at any other time if the cause were not present. Consider the April rains. They make it more probable that the forest fire occurs in June. They don't make the forest fire more probable per se. If we compare the p(forest fire in June) just before June—there having been April rains—with the p(forest fire in May), say, just before May—there having been no April rains—we find that its not the case that p(forest fire in June) >> p(forest fire in May). Nor are matters helped if we put events in Σ which would ensure that, even if the April rains did not occur, the forest fire would not occur before June—for instance, the non-actual events comprising the various ways in which the forest might have caught alight before June—so drastically lowering p(forest fire in May). If we did that, then the April rains would not even raise the chance of the forest fire occurring in June. The absence of the other events would already have ensured that the forest would not have caught fire before. The April rain will not be needed to guarantee the forest is not burnt down and could catch light in June. Similar considerations explain why the doctor's treatment is not a cause of the severe phase of illness.

We can test the strength of the proposal by considering the way it handles the alleged asymmetry between hasteners and delayers identified (although now repudiated) by Jonathan Bennett.

> Hasteners are usually causes of what they hasten; delayers are not usually causes of what they delay. (Bennett 1987, p. 375)

I have given some illustrations of delayers which are not causes. Here are a couple of cases of hastening mentioned by Penelope Mackie (1992, pp. 483–4). An invitation to give a paper in November hastened its completion which would otherwise have been in December. In this case, it seems intuitive to say that the invitation caused the completion of the paper. Smith has heart disease and would have had a heart attack on Saturday (when she ran a marathon). However, she has a row with her employer on the Wednesday beforehand and has a heart attack then. It seems intuitive to say that the row caused the heart attack. In both these cases, the account

[12] This proposal was developed with Murali Ramachandran to deal with the problem of late preemption. He deserves credit for the final formulation.

seems to get the right results. In the case of the paper, we would put in Σ the events that would have brought about the completion in December. So, in effect, what we would be comparing are the following two conditionals

(a1) If I were to receive the invitation without any of the events in Σ occurring, it would be the case that p(the paper is finished at *t*) is at least *x* (where the probability is assessed at the appropriate point in November just before *t*).

(a2) If neither were I to receive the invitation nor were any of the events in Σ to occur in December, it would be the case that p(the paper is finished at *t*) is at most *y* (where the probability is assessed at the appropriate point in December just before *t*).

It is clear that relative to the absence of these events in December, the invitation would raise the chance of completion in November to any other time, with December being the most likely time at which *y* would be largest. Similarly, we would put in Σ the event of running the Marathon (and doubtless many other events which may also give rise to that heart attack but let's keep things simple). The row would then raise the chance of the heart attack on Wednesday to its chance at any other time. Here the two conditionals would be

(b1) If she were to have a row with her employer and were not to run the Marathon, it would be the case that p(she has a heart attack at *t*) is at least *x* (where the probability is assessed at the appropriate point on Wednesday).

(b2) If she were neither to have a row with her employer nor to run the Marathon, it would be the case that p(she has a heart attack at t) is at most *y* (where the probability is assessed at the appropriate point on Sunday (say) just before *t*).

Again *x* would be very much greater than *y*. This is not what we found in the case of the April rains. When we put possible events in Σ which ensured that the forest did not burn down in June, the April rains did not raise the probability of the fire at all.

Although I put certain events in Σ which occur after the actual time of the effect, I do not hold that this makes any difference to the chance of the effect at that time. Rather, the chance it influences is the chance of the effect at the times at which it would occur given the absence of the cause. Moreover, the chance that I am assessing is the chance of an event occurring *at a certain time* assessed just before that time. For instance, consider a time just before the woman runs the Marathon. The chance of her having a heart attack some time after that time may be quite high (given that it is not settled that she is not running the Marathon). However, the chance of her having a heart attack at precisely that time before the Marathon is run is low regardless of whether it is settled that she will run it.

The proposal seems to fit nicely with Mackie's diagnosis of the difference between hasteners which are thought to be causes and delayers which are not. She suggests that *hasteners* generally *prevent* events from occurring later *by bringing them about earlier*. On the other hand, delayers bring about events later that would otherwise have occurred earlier *by preventing them from occurring earlier* (Mackie 1992, pp. 493–5). By insisting that a cause must—given that certain other events fail to occur—substantially raise the probability of an event at a time over other times, the proposal rules out *mere* time-switchers: events which themselves do nothing to bring about the effect but just determine the time at which it is brought about. The proposal rules in time-switchers that switch by bringing things about, that is, hasteners.

The proposal also manages to explain why certain delayers are causes. A familiar example is the fatal antidote to a poison. The subject would have died if he or she had not taken the antidote to the poison just ingested. The subject dies anyway, later, from a reaction to the antidote. Intuitively, the antidote both delays and causes the death. My proposal gets that result. Put the ingestion of the poison in Σ. Then the antidote raises the chance of the death later over the chance of the death at any other time due to the adverse reaction (Mackie 1992, pp. 485, 495–6). On the other hand, my proposal fails to accord with the intuitions of those who take some hasteners not to be causes. For instance, suppose I receive a subpoena requiring my attendance at the Old Bailey in September—just when I was planning to have my holiday in Paris. Suppose one puts in Σ events necessary for the occurrence of the holiday in September. Then receiving the subpoena makes the chance of the holiday in August very much higher than it would be at any other time. So my account claims that receiving a subpoena is a cause. I think that this is the right result in spite of some people's intuitions to the contrary.[13] To undermine the intuition that this is not correct, it is helpful to note two things. First, while a subpoena to give evidence in September prevents one from going on holiday then, *receiving the subpoena* actually brings about the holiday in August prior to the prevention. It doesn't fall under Mackie's "brings about by preventing" formula (Mackie 1992, pp. 494–5). Receiving the subpoena brings about the holiday because one believes one *will be* prevented from going on holiday in September. The prevention hasn't occurred yet. The subpoena may be withdrawn. Second, I think the residual intuition that the subpoena is not a cause is explained by the thought that, *in the circumstances*, the subpoena makes the holiday less likely. We have already seen, though, that this is not always indicative of whether something is a cause. For instance, in the early preemption case, *b*'s firing lowered the chance of

[13] I am grateful to an anonymous referee for pressing on me that I should say this.

e's firing in the circumstances. However, excluding the incomplete *a*-chain made *b*'s firing into a chance raiser. The subpoena case is comparable. Receiving the subpoena blocks the causal chain leading to a holiday in September and initiates the chain leading to the holiday in August. I am not just claiming here that receiving the subpoena caused the holiday to occur *in August*. I am claiming that receiving the subpoena in fact caused the holiday—just as *b*'s firing caused *e*'s firing not *e*'s firings, at a time. Events in a chain leading to the holiday in September could not have done this because the chain was preempted.

One small refinement is needed to complete the treatment. Suppose that the firing of a neuron, *a*, satisfies the first clause of the account by raising the chance of another neuron, *e*, firing at *t* over any other time. To simplify matters, suppose that the *a*–*e* connection is direct. There are no intervening events. Finally, suppose that *e* has some background chance of firing anyway and does so at time $t + 1$, a time at which the firing of *a* does nothing to make the firing of *e* more likely. Intuitively, the firing of *a* is not a cause of *e* firing but it seems that the firing of *a* would satisfy both clauses of my account. So we need one more clause:

(III) e_2 occurs at one of the times for which $p(e_2 \text{ at } t) \geq x >> y$.[14]

It is time to return to the case of probabilistic late preemption—Figure 2—fortified by this success. First consider whether the firing of *b* comes out as a cause. Let Σ include the firing of *d*. Then the chance of *e* firing at the time it did given that *b* fired would be very much greater than any of the chances that *e* has of firing at a time if *b* did not fire. So the firing of *b* satisfies clause (I) of the account of causation—the probabilistic Σ-dependence clause. It also satisfies the second clause—the "actual events" clause—because there are no non-actual events on the *b*-chain. Consider now the firing of *a*. Given that *e* has a background chance of firing anyway, it will still occur prior to *d*'s firing. If we consider the chance of *e* firing later—intuitively at the time it would have been brought about by the *a*–*e* process—then it is 0 whether or not *a*'s firing occurred. The very same event of *e* firing can't occur at two times. So the chance of *e* firing later—assessed just before *e* would then have fired—would already take into account the fact that *e* had fired earlier. Hence *e*'s firing can't probabilistically Σ-depend upon *a*'s firing, *a*'s firing fails clause(I) of my account. What if *e* does not have a background chance of firing anyway? Then the difference between the firing of *a* and the firing of *b* is revealed by their respective times *t* at which p(*e* fires at *t*) $x >> y$. In the case of the firing of *b*, one of the times is the time, t_o, at which *e* actually occurred, whereas, in the case of the firing of *a*, this is not so. P(*e* fires at t_o) is not

[14] This clause is also needed to deal with cases of frustration (see Noordhof 1998a, pp. 458–9).

raised by the presence or absence of the firing of *a*. If *e* fired as a result of the *a*–*e* process, it would have been later. So the firing of *b* passes and the firing of *a* fails clause (III).

The three clauses of the account can now be understood to work as follows. Clause (I) explains what it is for e_1 to be a cause of e_2 and not just a cause of e_2's occurrence at *t*. Clause (III) makes sure that e_1 in fact stands in the appropriate relationship to e_2 for the time at which e_2 occurred. Clause (II), together with the decision to assess p(*e* at *t*) just before t, specifies what would have to be the case at the time at which e_2 occurs for the causal chain between e_1 and e_2 to be complete.

Suppose that there is a certain kind of compound *X* which, when initially formed, is highly unstable. After a certain critical time period—five seconds say—it then becomes relatively stable and is only likely to break down under bombardment by subatomic *Y*-particles. Its chance of breaking down during a 1 second interval after it is bombarded *outside* the critical period is 0.6 whereas its chance of breaking down during a one second interval *in* the critical period—whether or not it is bombarded—is 0.7. If *X* is bombarded during the critical period, the chance of its breakdown is still around 0.7 because, let us imagine, the bombardment interacts with what makes the compound more stable which, during the critical period, is not present. Suppose that a bombardment takes place after the critical period. Intuitively, we might want to say that the bombardment caused the breakdown of the compound. However, it does not look as if my proposal can get that verdict. The probability of there being a breakdown if there is a bombardment is at least 0.6. But the probability of there being a breakdown in the critical period in the absence of the bombardment is 0.7. So the bombardment does not pass the "probabilistic Σ-dependence" clause of my proposal.

On the assumption that the very same breakdown can occur both during the critical period due to instability and after the critical period due to bombardment, I claim that we should deny that the bombardment is a cause of the breakdown. Instead, we should say that it is a cause of the breakdown occurring *during a certain time period*—the period after the critical period. We might capture this notion as follows.

> For any actual, distinct events e_1 and e_2, e_1 causes e_2 during time period *T* iff there is a (possibly empty) set of possible events Σ such that
>
> (I) e_2 is probabilistically Σ-*T*-dependent on e_1, and
>
> (II)′ For any superset of Σ, Σ*, (where $\Sigma \subseteq \Sigma^*$), if e_2 probabilistically Σ*-*T*-depends upon e_1, then every event upon which e_2 probabilistically Σ*-*T*-depends is an actual event, and
>
> (III) e_2 occurs at one of the times for which p(e_2 at *t*) $\geq x >> y$

where

e_2 probabilistically Σ-*T*-depends upon e_1 if and only if

(1) If e_1 were to occur without any of the events in Σ, then for some time *t* in *T*, it would be the case that, just before *t*, p(e_2 at *t*) ≥ *x*

(2) If neither e_1 nor any of the events in Σ were to occur, then for *any* time *t* in *T*, it would be case that, just before *t*, p(e_2 at *t*) ≤ *y*

(3) $x >> y$.

Here I have just limited the account to a certain time period—that represented by *T*. Causes of e_2 during a time period *T* make the chance of e_2 very much greater than the maximal background chance of e_2 during that time period. Causes of e_2 *per se* make the chance of e_2 very much greater than its maximal background chance at any time.

A recommendation of this approach to what I take to be a borderline case is that it keeps the distinction between causing an event to occur and causing it to occur at a time while not involving itself in an unmotivated stipulation. If we claim that e_1 may be a cause of e_2 because it raises its chance in the required way merely *during a certain time period*, then we need to settle how small the time period could be for us still to have a genuine case of causing e_2 as opposed to causing e_2 to occur at a time or during a certain time period. There seems no obvious resolution of this matter. By contrast, my proposal is that the cause of an event will make the chance of it occurring at one time very much greater than the chance it has of occurring at any other time (given the events in Σ don't occur). No arbitrary resolution is required.

4. Catalysts and anti-catalysts

One residual problem area for the type of approach defended here—appealing to times—are cases of preemption in which one of the processes is a catalyst or anti-catalyst of another process. Here is an example of what I have in mind.

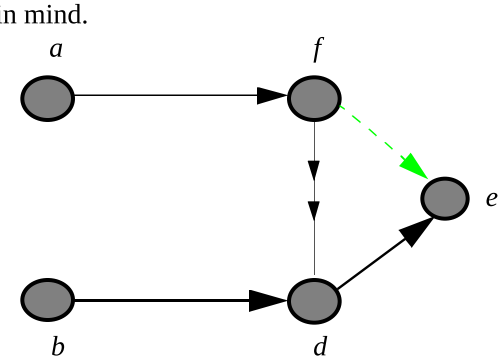

Figure 4

The crucial feature is the inhibitory axon between *d*'s firing and *f*'s firing. It does not stop *f* firing but makes it fire later (it acts as an anti-catalyst). As a result, *e* fires at *t* (say) because of the *b*-chain and *f*'s firing fails to raise the chance of *e*'s firing at *t* which it would have done if *f*'s firing had not been slowed down. If the *b*-chain had not caused *e* to fire at *t*, *f*'s slowed firing would have still raised the chance of *e* firing later. My proposal does not have a problem with getting the verdict that *b*'s firing is a cause of *e*'s firing. Just let Σ = the firing of *a*. The problem is with whether the firing of *a* is a cause. If we put either *b*'s or *d*'s firing in Σ, it would seem that the firing of *a* would satisfy clauses (I) to (III).

The new feature introduced by this type of case is that the appeal to Σ-sets may change the time at which a certain process will bring about an effect e_2. What we need to do is make sure that not only is the causal chain between e_1 and e_2 complete (clause (II) et al.) and that the right relationship holds between e_1 and e_2 at the time at which e_2 occurred (clauses (I) and (III)) but also that e_1 is related to e_2 occurring at the time it did *in the actual circumstances* and not merely in some possible circumstances. The case considered in Figure 4 dramatizes this point.

The following clause is aimed to capture what we need.

(IV) e_2 probabilistically *A*-time depends upon e_1.

This is defined as probabilistic Σ-time-dependence was defined above, with *A* in place of Σ. It is specifically defined for t_o, the actual time e_2 occurred.

> e_2 probabilistically *A*-time-depends upon e_1 if and only if there is a (possibly empty) set of possible events *A* such that
>
> (1) If e_1 were to occur without any of the events in *A*, then it would be the case that $p(e_2 \text{ at } t_o) \geq x$
>
> (2) If neither e_1 nor any of the events in *A* were to occur, then it would be the case that $p(e_2 \text{ at } t_o) \leq y$
>
> (3) $x >> y$.

It might be wondered why I appeal to probabilistic *A*-time-dependence when I rejected something of a similar form when initially trying to characterize the kind of probabilistic dependence that must hold between two events if they are to be related as cause and effect. The answer is that, at this point, all I am testing for is whether a particular event has influenced the probability of another event occurring at the actual time it did; I am not trying to capture the causal relationship in general.

The reason for appealing to another set, *A*, is that the members of Σ won't necessarily be appropriate for *A* as we saw in the case described above. The aim of this clause is to test whether, *as things are*, e_1 has an influence upon e_2 at the time it occurred—not as things might be in circumstances in which the events in Σ don't occur. Doubtless it will seem that I am just introducing a new clause for which the same problem arises.

But I am not because—in contrast with Σ—there must be constraints on what can be a member of *A*. Specifically

No event e_i can be put in *A* if both

(a) If it is a member of *A*, then <e_1, e_2> satisfies (IV)

(b) If it is not a member of *A* and we replace (IV)(1) and (2) with

(1*) If e_1 and e_i were to occur, *with none of the events in A occurring, nor e_i satisfying any of (I) to (III) regarding e_2*, then it would be the case that p(e_2 at t_o) ≥ *x*

(2*) If e_i were to occur with neither e_1 *nor any of the events in A occurring, nor e_i satisfying any of (I) to (III) regarding e_2*, then it would be the case that p(e_2 at t_o) ≤ *y*

then <e_1, e_2> do not satisfy (IV).

In effect, what this constraint does is make sure that the same problem does not arise by insisting that the only events that may be members of A are events whose absence leave untouched the relationship between the candidate cause, e_1 and the p(e_2 at the time it actually occurred). The clause works like this. We are supposed to consider what would be the case if e_i were not to satisfy clauses (I) to (III) of my account. This way we can focus on the influence of e_i on the e_1–e_2 connection independently of any role that e_i might play in raising the chance of e_2. In other words, we can focus entirely on its catalytic or anti-catalytic features. If e_i does not satisfy clauses (I) to (III) in the actual world, little would have to change. If e_i does satisfy these clauses, then the laws and particular circumstances would have to change just enough so that the clauses cease to hold of e_i and e_2. Then, if we get a different verdict about the relationship between *x* and *y* for (1*) and (2*) than we do for (1) and (2) above, we know that e_i has an influence on the relationship between e_1 and p(e_2 at t_o) in addition to its independent influence on the chance of e_2. So we cannot safely put e_i in *A*.

The thinking behind the constraint is easier to see when we apply it to the case which troubled us (and a related case). First, regarding Figure 4, it allows us to pronounce the firing of *b* to be a cause. All we have to do is let *A* = {*f* fires). The firing of *f* would not be ruled out from being a member of *A* since if we consider the relevant conditionals

(a1) If the firing of *b* and the firing of *f* were to occur without *the firing of f satisfying any of (I) to (III) regarding the firing of e*, then it would be the case that p(the firing of *e* at t_o) ≥ *x*

(a2) If the firing of *f* were to occur with neither the firing of *b* nor *the firing of f satisfying any of (I) to (III) regarding the firing of e*, then it would be the case that p(the firing of *e* at t_o) ≤ *y*.

x would still be very much greater than *y*. The relationship would be unchanged. This is no surprise since it is clear that *f*'s firing has no influence on the causal chain from *b* firing to *e* firing. By contrast, the firing of *d* would be ruled out from being a member of *A* for the firing of *a*. The appropriate conditionals this time would be

(b1) If the firing of *a* and the firing of *d* were to occur without *the firing of d satisfying any of (I) to (III) regarding the firing of e*, then it would be the case that p(the firing of *e* at t_o) ≥ *x*

(b2) If the firing of *d* were to occur with neither the firing of *a* nor *the firing of d satisfying any of (I) to (III) regarding the firing of e*, then it would be the case that p(the firing of *e* at t_o) ≤ *y*.

If *d*'s firing occurred without satisfying (I) to (III), then the *b–d* chain would not raise the chance of the firing of *e* at time t_o. Neither would *a*'s firing. The firing of *d* would still occur to slow the *a–e* chain down so that *a*'s firing would not raise the chance of e_2 at t_o. So it is not the case that *x* >> *y*. The verdict has changed. So *d*'s firing cannot be a member of *A* and, as a result, the firing of *a* can't pass (IV).

To confirm the application of this clause, consider the figure below.

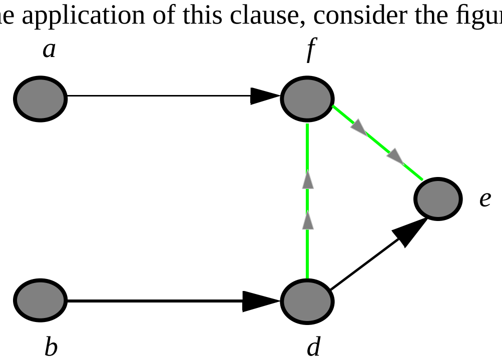

Figure 5

Here *d*'s firing is a catalyst speeding up the *a–e* process. If *d* had not fired, then the top process would have been slower and failed to complete by the time that *e*'s firing occurred. As things are, it completed. For the same reasons as those above, we may conclude that *b*'s firing is a cause—for instance by letting *f*'s firing be a member of *A*. However, matters are more complicated in the case of *a*'s firing. If the *b–e* chain is indeterministic, then the firing of *a* may be a cause if the *a–e* chain raises the chance of *e* firing significantly over the chance it has as a result of the firing of *b*. Suppose the *a–e* chain doesn't raise the chance significantly. Then the proposal can't establish that *a*'s firing is a cause. Letting *A* = {*d* firing} will do no good because, with the catalyst gone, the firing of *a* will fail clause (IV) of the account. It won't raise the chance of *e* at t_o. (Obviously the firing of *d* can be a member of *A* since it passes the membership rules by failing clause (a), the clause requiring that if the firing of *d* is a member of *A*, *a*'s firing, *e*'s firing must satisfy clause (IV) of my account.) On the other hand, if the *b–e* chain is deterministic, then the firing of *a* will not be a cause of *e* firing since it can't raise the chance of *e* firing. It seems to me that these verdicts are plausible. When the firing of *b* not only raises the chance of *e* but also acts as a catalyst for another process, the other process only has independent causal credentials if it very much raises the chance of *e* firing. Of course, this does not stop us claiming that the firing of *a and b* are a collective cause.

I think I have now taken the approach to the outskirts of firm intuition. However, if you are convinced that the firing of *a* is a cause, then let me briefly show you how my proposal may be developed to obtain that result. This should demonstrate the resilience of the approach and undermine scepticism about its overall line of thought. My proposal would need to be altered in two ways. First, clause (IV) should read

> e_2 probabilistically *A*-time depends upon e_1, or e_2 probabilistically $A^{\#}$-time depends upon e_1

(where probabilistic-$A^{\#}$-time dependence is defined by substituting (1*) and (2*) for (1) and (2) in the definition of probabilistic *A*-time dependence). The relevant conditionals in the present case would then be

(c1) If the firing of *a* and the firing of *d* were to occur without *the firing of d satisfying any of (I) to (III) regarding the firing of e*, then it would be the case that p(the firing of *e* at t_o) $\geq x$

(c2) If the firing of *d* occurred with neither the firing of *a* nor *the firing of d satisfying any of (I) to (III) regarding the firing of e*, then it would be the case that p(the firing of *e* at t_o) $\leq y$.

Since $x >> y$, the firing of *a* would pass the revised clause (IV).

Second, the overall account would have to be adjusted so that e_1 and e_2 must satisfy either (III) or (V), where (V) should be consulted if e_1 and e_2 pass clause (IV) only by passing the probabilistic $A^{\#}$-time dependence component. Let the *causal shell of an event x* be *x* satisfying clauses (I) to (III), let $\Sigma^{\#}$ be Σ with one or more events dropped from the Σ-set and replaced by their causal shells, let the events dropped be the *replaced events*. Finally, let

> e_2 probabilistically $\Sigma^{\#}$-depends on e_1 if and only if

(1) If e_1 and the replaced events were to occur without any of the events in $\Sigma^{\#}$ or the causal shells, then for some time *t*, it would be the case that, just before *t*, p(e_2 at *t*) $\geq a$

(2) If the replaced events were to occur, but neither e_1 nor any of the events in $\Sigma^{\#}$, nor their causal shells, then for any time t, it would be the case that, just before *t*, p(e_2 at *t*) $\leq b$

(3) $a >> b$.

Then (V) should read:

> e_2 probabilistically $\Sigma^{\#}$-depends on e_1 and e_2 occurs at one of the times for which p(e_2 at *t*) $\geq a >> b$.

The reason for this second adjustment is roughly as before. It just acknowledges the fact that we are now appealing to causal shells as well. We need make sure that e_1 raises the chance of e_2 in the appropriate way at the time at which e_2 occurred.

5. A few defensive comments

It is time to draw the various strands of the account together and consider what it promises to provide. My proposal is that

For any actual, distinct events e_1 and e_2, e_1 causes e_2 (if and) only if there is a (possibly empty) set of possible events Σ such that

(I) e_2 is probabilistically Σ-dependent on e_1, and

(II)′ For any superset of Σ, Σ*, (where $\Sigma \subseteq \Sigma^*$), if e_2 probabilistically Σ*-depends upon e_1, then every event upon which e_2 probabilistically Σ*-depends is an actual event

(III) e_2 occurs at one of times for which $p(e_2 \text{ at } t) \geq x >> y$

(IV) e_2 probabilistically *A*-time depends upon e_1.

Although this formulation might appear complex, the final conception of a cause it articulates is relatively simple.

> A cause, e_1, is something which (independantly of its competitors) both makes the chance of an effect, e_2, very much greater than its maximal background chance (clause (I)) *and* actually influences the probability of the effect in this way at the time at which the effect occurred (clauses (III) and (IV)) via a complete causal chain (clause (II) and the way in which probabilities are assessed).

To keep this paper within manageable proportions, I don't want explicitly to advance an account of causal asymmetry here. That is why I have couched the final account as a necessary condition with brackets around the sufficiency claim. It awaits to be supplemented by an account of causal asymmetry (and that alone—I am prepared to say). However, I think it can be shown how Lewis's approach to asymmetry could be made to fit with the above account—which is not to say his account is without its own problems (see Vihvelin 1995, pp. 565–75; Hausman 1996, pp. 60–1). So, as far as the counterfactual approach is concerned, my proposal leaves things more or less undisturbed. I will briefly indicate why this is so.

Lewis's approach to causal asymmetry starts with the claim that backtracking counterfactuals—counterfactuals with effects mentioned in the antecedent and causes in the consequent—are not, in general, true. Thus, although it might be true that

> If they hadn't heard the ambulance siren, they wouldn't have moved out of the way,

Lewis denies that it follows that

> If they hadn't moved out of the way, then they wouldn't have heard the ambulance siren.

Instead, he suggests that they would have heard the ambulance siren but that other things would have occurred to stop them from moving out of the way—for instance, the road was not clear, they froze, and so on. He thinks that this is the right thing to say because possible worlds in which we envisage a putative cause to be absent (as well) are going to be worlds which are even less similar to our own than worlds in which merely the effect is absent. A cause has many other consequences in our world. In order to retain these consequences in the absence of the cause, we would have to allow a huge number of departures from the laws of our world. On the other hand, if we did not retain these consequences, then the world would be very different in particular matters of fact up to the time of the effect. Either way, the two most important standards of similarity between worlds would be infringed (Lewis 1979, pp. 47–8). So Lewis suggests that the closest worlds are ones in which, even if the effect had not occurred, the putative cause would still have occurred. The only violation would be in the laws linking cause and effect. Causal asymmetry is assimilated to counterfactual asymmetry, which is in turn elucidated in the way indicated.

Obviously, the clauses of my account, though they contain counterfactuals—are a step on from the simple counterfactuals for which Lewis denies the truth of their backtracking cousins. We need to be sure that the modifications which my account introduces does not allow that effects are causes of their causes. Does the account hold if we swap effect e_2 for cause e_1 in the clauses above?

By appealing to the notion of a Σ-set of events which fail to occur, it might at first be thought that I lose the asymmetry upon which Lewis's account rests. Suppose that the event of people hearing the ambulance siren had lots of consequences apart from their getting out of the way. They remarked upon it to their friends in the car, they turned round to look, they sweated slightly nervously and so on. All these can be put in the Σ set. Then it is surely true that if all these things had not occurred, p(people hearing the ambulance siren) would be very low. So it seems that I must conclude that their not getting out of the way is a cause of their hearing the ambulance siren.[15] However, that is to ignore the other component

[15] This is a problem which afflicts Ramachandran's account (Ramachandran 1997, see Noordhof 1998a, pp. 460–2). He cannot give the reply I offer because he appeals to might-conditionals rather than would-conditionals to characterize his dependency.

of my account of probabilistic Σ-dependence. If they were to move out of the way but were not to remark upon it to their friends in the car, nor turn round to look, nor sweat slightly nervously, … then p(they would have heard the ambulance siren) would also be very low. If all the other consequences of a particular cause did not happen and yet the putative effect did, it would be less of a departure from the actual world if we envisaged that the effect was brought about by other means rather than that all the connections between the cause and the other consequences failed to hold. So we still have our asymmetry. If the case is deterministic, this is particularly conspicuous. If they moved out of the way but none of the other consequences occurred, the probability of them hearing an ambulance siren would be 0. Effects will fail clause (I) of the account.[16]

Obviously, there are many other matters to discuss in order to ascertain whether the proposal defended here is correct, for example: cases of overdetermination, causation by fragile events, issues relating to the individuation of events, and the like. It would not be appropriate to go into these concerns here. My aim has been merely to respond to a particularly pressing problem for the counterfactual theory rather than give a complete defence. Nevertheless, I do not believe that they raise any new difficulties which merit adding clauses to the account. Given my overall aim, it is possible for me to duck more general objections to the counterfactual approach. One could view my discussion as just showing that, if a counterfactual approach is prima facie viable, then there are no particular reasons of substance for surrendering this approach because of problems relating to probabilistic causation and preemption. However, general objections to an approach have a habit of undermining interest in the details of a particular theory, so I think I do need to sketch a response to a line of scepticism that has been expressed recently.

A natural worry to have about a counterfactual theory of causation is that in arriving at our judgements concerning what would happen if such and such were the case, we are explicitly appealing to causal considerations. Daniel Hausman criticises Lewis's idea that in order to ascertain whether a counterfactual is true, we should envisage the world fixed up to a moment or so before the contrary-to-fact antecedent of the conditional, imagine a divergence of the laws of the world from the laws of our world (a "miracle") in order to bring about the circumstances envisaged in the

[16] Another source of the asymmetry is that there is more than one event—say *a*, *b*, and *c*—which counts as a cause of *e*. So if *e* were not to occur, either a or b or c would not have occurred. But that means that there is no backtracking conditional of the form "if *e* were not to occur, then *a* (say) would not occur" since there would be as close worlds as the not-*a* worlds in which *b* did not occur (Hausman 1996, pp. 58–61). For an alternative account of the asymmetry in terms of the present approach, see Noordhof (1998a, p. 462, fn. 3); for criticism of it see Ramachandran (1998, p. 469, fn. 7).

antecedent, and roll on the world from there given the laws which hold at that world. Instead, he suggests, the predictions made on the basis of counterfactual reasoning will be the result of identifying the *causal* background to the conditions mentioned in the antecedent. For instance, if we ask what would happen if a steam pipe burst in a nuclear reactor we need to consider the way in which this occurred: as a result of it being faulty? an earthquake? the falling of a girder? sabotage? or too great a pressure? As Hausman notes

> If the pipe burst because the pressure was too great, and the pressure was too great because the reactor was going out of control, then the consequences of the pipe bursting may be different than if it were caused by corrosion, a faulty weld or a terrorist's bomb. (Hausman 1996, p. 65)

We backtrack in order to think about potential causes of a cause, and then extrapolate forward to try to identify what would be the consequences given these background conditions. The role that miracles play is further back in the causal history when the conditions are not thought to matter to the effects under consideration (Hausman 1996, p. 69).

Hausman concludes from this that

> one must give up any hope of providing a counterfactual theory of causality, since part of one's basis for judging the similarity among possible worlds and the truth of counterfactuals would be explicitly causal. (Hausman 1996, p. 70)

However, I do not think that Hausman has managed to establish that we should give up all hope. It rather depends upon what we are aiming to do. The aim is not to provide an account of causation without appealing to our primitive grasp on the concept of cause. In formulating any analysis of a concept, appeal to our primitive grasp of it will be necessary. More importantly, the aim is not to give an analysis of causation in terms of something—counterfactuals—whose truth conditions one could grasp without any prior ability to grasp the truth conditions of our causal judgements. For instance, Lewis has quite happily acknowledged that there is a close connection between the counterfactuals we are inclined to assert—the worlds that we judge are close—and our causal knowledge (e.g. Lewis 1979, pp. 47–9, Lewis 1986, pp. 66, 211). This has not diminished his enthusiasm for a counterfactual theory.

In the hands of Lewis, the aim of a counterfactual analysis of causality is to establish that causal relations supervene upon properties and relations other than necessary connections. If the truth conditions of causal judgements can be cashed out in terms of arrangements of particular matters of fact and a specified similarity ordering of possible worlds—worlds within which there are no necessary connections—then he has succeeded in his aim. The notion of similarity between worlds has nary a hint of necessity to it. To defeat his programme, it would have to be established that causal facts do not supervene on these other things. To do this, it is not enough to note that one's counterfactual predictions about what would happen if a pipe were to burst depend upon whether this is through corrosion or the accumulation of pressure. These two hypotheses would correspond to differences in particular matters of fact which could be taken to be salient in determining the similarity ordering of worlds. There is nothing wrong with taking prior causal judgements as providing additional information about the similarities of particular matters of fact which matter. To threaten Lewis's approach, what one has to establish is that these differences in particular matters of fact (and the laws which govern them) by themselves would not determine the appropriate similarity ordering of worlds. One would have to advert to sui generis causal differences. I do not think that Hausman has established this stronger claim.

If I am right, then Lewis's aim would be advanced by the proposal I have developed here. As far as I can see, there is nothing I have introduced which threatens it. However, I would be quite satisfied merely to have achieved something weaker. We make causal judgements, assert counterfactuals and note probabilistic dependencies. Our understanding of these activities will be advanced if we note the connections between them. Our understanding will also be advanced if we appreciate the range of notions which are connected to our understanding of causality. If the theory put forward here is on the right lines, I have shown that our understanding of causation is connected to our understanding of counterfactuals, probabilities and nothing else. I would prefer to leave it to others to assess claims of priority. [17]

Department of Philosophy
University of Nottingham
University Park
Nottingham NG7 2RD
UK

PAUL NOORDHOF

[17] At two places in this paper, I have remarked that Murali Ramachandran deserves credit for the final formulation of a component of the theory. Let me acknowledge once more the important role he played in the development of the ideas put forward here. I very much doubt whether I would have arrived at them alone. Our exchanges during the summer of 1997 remain the most fruitful collaborative work I have experienced. Perhaps I should also note that he does not accept my theory and is developing his own! I would also like to thank Jonardon Ganeri whose earlier collaboration and more recent discussion has also been very helpful and Michael Clark who raised a whole host of fascinating issues and cases which deserve more attention than I could give them here. Lastly, but certainly not least, I would like to thank two anonymous referees. One of them gave me among the most detailed, constructive and engaging comments that I have been lucky enough to receive from a journal referee. Both saved me from a number of serious errors.

Paul.Noordhof@nottingham.ac.uk

REFERENCES

Bennett, Jonathan 1987: "Event Causation: The Counterfactual Analysis", in James E. Tomberlin (ed.), *Philosophical Perspectives I*. Atascadero, California: Ridgeview Publishing Company, pp. 367–86.

¾¾ 1988: *Events and their Names*. Oxford: Oxford University Press.

Byrne, Alex and Hall, Ned 1998: "Against the PCA-Analysis". *Analysis*, 58, pp. 38–44.

Ganeri, Jonardon, Noordhof, Paul and Ramachandran, Murali 1996: "Counterfactuals and Preemptive Causation". *Analysis*, 56, pp. 216–25.

¾¾ 1998: "For A (Revised) PCA Analysis". *Analysis*, 58, pp. 45–7.

Hausman, Daniel M. 1996: "Causation and Counterfactual Dependence Reconsidered". *Noûs*, 30, pp. 55–74.

Honderich, Ted 1988: *A Theory of Determinism*. Oxford: Oxford University Press.

Lewis, David 1973: "Causation", in Lewis 1986, pp. 159–72. Originally published in 1973 in *Journal of Philosophy*.

¾¾ 1979: "Counterfactual Dependence and Time's Arrow", Lewis 1986, pp. 32–66. Originally published in 1979 in *Noûs*.

¾¾ 1986: *Philosophical Papers*, Volume 2. Oxford: Oxford University Press.

Mackie, Penelope 1992: "Causing, Delaying, and Hastening: Do Rains Cause Fires?". *Mind*, 101, pp. 483–500.

McDermott, Michael 1995: "Redundant Causation". *The British Journal for the Philosophy of Science*, 46, pp. 323–44.

Mellor, D. H. 1995: *The Facts of Causation*. London: Routledge.

Menzies, Peter 1989: "Probabilistic Causation and Causal Processes: A Critique of Lewis". *Philosophy of Science*, 56, pp. 642–63.

¾¾ 1996: "Probabilistic Causation and the Pre-emption Problem". *Mind*, 105, pp. 85–116.

Noordhof, Paul 1998a: "Problems for the M-Set Analysis of Causation". *Mind*, 107, pp. 457–63.

¾¾ 1998b: "Critical Notice of D. H. Mellor, *The Facts of Causation*". *Mind*, 107, pp. 855–75.

Ramachandran, Murali 1997: "A Counterfactual Analysis of Causation". *Mind*, 106, pp. 263–77.

¾¾ 1998: "The M-Set Analysis of Causation: Objections and Responses". *Mind*, 107, pp. 465–71.

Vihvelin, Kadri 1995: "Causes, Effects and Counterfactual Dependence". *Australasian Journal of Philosophy*, 73, pp. 560–75.