

CHANCE-CHANGING CAUSAL PROCESSES

Helen Beebee

Please do not cite this version. The published version is: ‘Chance-Changing Causal Processes’ in P. Dowe and P. Noordhof, eds., *Cause and Chance* (London: Routledge, 2004), 39-57

1. Causation, Chance increase, and causal processes

Consider the following two *prima facie* plausible claims about causation:

- (CP) Causes and effects must always be connected to each other via a causal process.
- (IC) Causes must always raise the chances of their effects.

For the purposes of this paper, I shall not attempt anything like a formal definition of a causal process, but will rest content with a common-or-garden, intuitive conception.¹ For example, there is a causal process between my hitting the white ball with the cue, the white hitting the black, and the black landing in the pocket. There is a causal process between my being in the presence of someone with flu, my getting flu, my gradual recovery (involving my antibodies fighting the infection and so on), and my return, a week later, to full health. There is no causal process between my writing these words, the neighbour’s dog barking, and the light in the next room being on. There is no causal process between my failure to go to the supermarket and my subsequent failure to cook dinner.² I shall also assume that (IC) is to be read counterfactually: *c* increases the chance of *e* if and only if, had *c* not occurred, the chance of *e* (just after the time at which *c* in fact occurred) would have been lower than it actually was.

As stated, each of (CP) and (CI) claims to represent a necessary, though not sufficient, condition for causation. Some stock examples suffice to show why neither the obtaining of a causal process nor increase in chance should (individually) be taken to be a sufficient condition for causation. First, it is generally agreed that there can be a causal process between *c* and *e* without its being true that *c* caused *e*; the example given above of my getting flu and my subsequent healthy state a week later will do. A structurally similar and well-known case, which I shall call the defoliant case, involves a plant being sprayed

¹ For attempts at a formal definition of a causal process, see for instance Dowe (1992) and Beebee (1997).

² This latter example is perhaps more controversial than the others are. One might, for instance, hold that there is a causal process of negative events – including, say, the emptiness of the fridge, my not chopping onions, and so on – linking my failure to go shopping to my failure to cook dinner (provided that an appropriate relation – counterfactual dependence, say – holds between each step). Anyone who holds the view that causal processes can obtain between absences, or between events and absences or absences and events, should note that I mean by ‘causal process’ the more common-or-garden kind – a kind that does not count negative events as participants in genuine causal processes.

with defoliant (*c*), recovering, and eventually being in full health again (*e*).³ There is a genuine causal process between *c* and *e*, but nobody, so far as I am aware, claims that *c* caused *e*. Second, there can be increase in chance without causation. Fred and Ted both want Jack dead. Fred poisons Jack's soup and Ted, unaware of Fred's act, poisons Jack's coffee. Suppose that each act increases the chance of Jack's death. Jack eats the soup but, feeling rather unwell, leaves the coffee – and dies later of poisoning. Ted's act raised the chance of Jack's death but did not cause it.

A simple diagnosis of the two kinds of case suggests a straightforward way of providing a sufficient condition for causation. In the flu and defoliant cases there is a causal process between *c* and *e*, but *c* lowers, rather than increases the chance of *e*: getting the flu decreases my chance of being fully healthy a week later, and being sprayed with defoliant decreases the plant's chance of being fully healthy six months later. In the poisoning example, on the other hand, while *c* increases the chance of *e* there is no causal process between the two: no chain of events or process links Ted's act with Jack's death, since Jack did not so much as go near the poisoned coffee. Hence, by taking the causal process condition and the chance increase condition to be *jointly* sufficient for causation, we rule out all the problem cases in one neat manoeuvre.

So far so good. Now, if only those conditions were also *necessary* for causation – that is, if only (CP) and (IC) were true – we would have ourselves necessary and sufficient conditions for causation.⁴

First, (CP). (CP) sounds plausible enough, but on closer inspection its truth is by no means obvious. For one thing, it rules out at least some cases of causation by absence, and also prevention (understood as the causing of the absence of an event). For another thing, it rules out causation at a temporal distance, since such causation would, by definition, involve the causing of one event by another without the aid of any process linking the two. (Strictly speaking, it also rules out direct, but not at-a-distance, causation. If time is quantised then, between two events that are as close to each other in time as it is possible to get, there cannot be any further events – so, strictly speaking, there can be no causal process between them. We could circumvent this problem by characterising a causal process as a process that is 'non-gappy' rather than as a sequence of events between which one can always interpolate further events that hook them together. But I shall leave such technical issues aside.) An adequate defence of (CP) is a big (some will doubtless think impossible) job; but for the purposes of this paper, I shall set such worries aside and assume that (CP) is true.⁵

What about (IC)? Well, there are some alleged counter-examples to it. One alleged counter-example is as follows: Sue pulls her golf drive (*c*), thereby lowering her chance of a hole-in-one. Fortunately, however – and against the odds – having hit a tree and bounced back onto the fairway, the ball lands in the hole (*e*). *c* lowered the chance of *e*,

³ The defoliant case is due to Nancy Cartwright (1979).

⁴ Assuming, of course, that we could provide satisfactory analyses of causal processes and counterfactual increase in chance – no easy task, and not one I shall attempt here.

⁵ I make a start on this job in my (forthcoming), where I defend the view that there is no causation by absence.

but one might still be inclined to say that c caused e .⁶ Another alleged counter-example runs as follows. An old lady – let’s call her Edna – is crossing the street, just as a bus comes hurtling towards her. I see the bus, and push Edna onto the pavement, out of the path of the bus (c). Unfortunately, however, I push her into the path of a falling brick. The brick is not falling from a great height, and Edna is wearing a sturdy hat; nonetheless, improbably, the brick hits her on the head and she dies. It seems right to say that my push caused Edna’s death, even though in pushing her I greatly reduced its chances.⁷

A standard move at this point is to claim that Edna’s actual death-by-brick is a different event to the event – the death-by-bus – that probably would have occurred had I not pushed Edna onto the pavement. This move restores chance increase, since had I not pushed Edna, her chance of dying the particular death she actually died (rather than some other death) would have been zero. In the current case, this move seems entirely plausible. There is no *a priori* reason to hold that whenever one attempts, but fails, to save a life, the actual death that occurs is the very same death that would have occurred in the absence of the attempt; the identity of the two deaths may hold in some cases but not in others. Compare, for example, the current case with a case where I push Edna but (say because Edna is very big and I am very weak) simply fail to move her out of the bus’s path. In the former case, the push initiates a causal process that is entirely unlike the causal process that would have continued in the absence of the push, and results in a death whose manner is entirely unlike the manner of the death that would (probably) occurred in the absence of the push. In the latter case, the death that actually occurs and the causal process that leads to it are (perhaps not precisely, but more or less) the same in manner as the causal process and death that would (probably) have occurred in the absence of the push. Plausibly it is only the latter case that the two deaths are identical. But in the latter case we are much less inclined to say that the push was a cause of Edna’s death.

Moves like this are not, unfortunately, always possible. Suppose that the actual brick-death and the non-actual bus-death are indeed two different deaths, but that their precise time and manner are not so different as to affect later consequences. Suppose, for example, that whichever death Edna were to die, her funeral would be conducted at the same time and place, and in the same way. Call the funeral event f . Then while c does not (if the brick-death and the bus-death are sufficiently different in manner to count as different deaths) lower the chance of e , it does lower the chance of f : had c not occurred, the chance of that very funeral’s happening just as it did would have been greater, since Edna’s chance of dying (as opposed to dying the particular death she actually died) would have been greater had I not pushed her. For the purposes of simplicity I shall assume that the death-by-bus and the death-by-brick are the very same event. However, everything I say about the relation between the push (c) and Edna’s death (e) might just as well be said about the relation between c and f , so readers who deny that the death-by-bus and the death-by-brick are the same event should substitute f for e in subsequent discussion of the case.

⁶ Dowe claims that the pulled drive is a cause of the hole-in-one, though others (including Hugh Mellor, whose example it is) disagree; see Mellor (1995), pp. 67-68.

⁷ See Dowe (2000), p. 71 for this example, henceforth called ‘the bus-brick case’.

If common sense intuitions about such cases are to be taken seriously (though I shall argue later that they need not), the chance-increase condition cannot be taken to be necessary for causation (which is to say, (IC) cannot be true), since in such cases *c* causes but lowers the chance of *e*.

Alleged cases of chance-decreasing causation create a problem not just for the prospects of providing necessary and sufficient conditions for causation in terms of causal processes plus chance increase; they also create a problem for what Phil Dowe calls the ‘Chance Changing Thesis’ (hereafter (CCT)):

(CCT) *c* promotes or tends to cause *e* when *c* raises the chance of *e*, and *c* causes *e* when *c* successfully promotes *e* (i.e. when *c* raises the chance of *e* and *e* occurs).

c hinders (or inhibits) *e* when *c* lowers the chance of *e*, and *c* prevents *e* when *c* successfully hinders *e*, i.e. when *c* lowers the chance of *e* and *e* fails to occur.⁸

(CCT) is an appealing thesis, and one that Dowe wants to uphold. However, since (CCT) implies (IC), counter-examples to (IC) are counter-examples to (CCT) too. He therefore adopts a strategy for dealing with the counter-examples that runs as follows. First, he diagnoses the problem cases – cases of chance-lowering causation – as cases where there is a ‘mixed path’ from *c* to *e*. Roughly speaking, the idea is this. In mixed path cases, *c* initiates two ‘processes’: one that could lead to – or tends to cause – *e* and one that could prevent, i.e. hinders, *e*. Since the hindering process is stronger than the promoting process, *c* lowers the chance of *e*.

Dowe’s solution to the problem of chance-lowering causes is to distinguish between (what I’ll call) the ‘all things considered’ chance of *e* (this is the chance of *e* that *c* lowers) and the chance of *e* ‘relative to a process’. He argues that if we abstract away from the hindering process and consider the chance of *e* relative just to the promoting process, then chance increase is restored: relative to this process, the chance of *e* in troublesome mixed path cases really is increased by *c*. Thus (CCT), and therewith (IC), is saved – once we interpret ‘chance’ as ‘chance relative to a process’ rather than the usual ‘all things considered’ kind.

For example, in the bus-brick case, pushing Edna out of the path of the bus and into the path of the brick initiates two different ‘processes’. On the one hand, the push initiates a genuine causal process (in the intuitive sense described earlier) which in fact culminates in Edna’s demise. On the other, the push initiates another ‘process’ that might (and indeed is intended to) save Edna’s life. (This ‘process’ involves Edna’s not being in front of the bus as it speeds along, not being hit by it, and so on. This is *not* a genuine causal process in the sense described earlier, since it is a sort of chain of non-events rather than a chain of events.) Unfortunately for Edna, this hindering ‘process’ is unsuccessful – it fails to prevent her death.

The chance-relativising thought is roughly this: abstract away from the (hindering) bus-avoiding ‘process’ and concentrate just on the (promoting) brick-process. Intuitively, just

⁸ See Dowe (2000), p. 69 (I have changed the wording slightly).

imagine that the bus was never there. In such an imagined scenario, the push is not a potential death-preventer, since had Edna remained in the middle of the road she would have been perfectly safe. The brick process in the imagined scenario, however, *is* still there. So in that scenario – which is to say, relative to the brick process – the push really does increase the chance of Edna’s death, since had I not pushed, Edna would have been very much less likely to die.

Like Dowe, I want to hang on to (IC). However, as I argue in section 2, his strategy for saving (IC) is unsuccessful. In section 3, I present a different conception of ‘hindrance’, according to which hindrance is a causal relation manifested by chance-lowering causal processes. I show how analyses of causation that do not count hindrance as a *bona fide* causal relation face a major problem, and argue that, with the appropriate notion of hindrance firmly in place, biting the bullet with respect to alleged cases of chance-lowering causation is a plausible strategy. The bullet-biting strategy simply denies that common-sense intuitions concerning the alleged counter-examples deserve to be taken seriously enough to undermine (IC), and I argue that the desire to believe in chance decreasing causes can be explained away. Finally, in section 4, I argue that the alleged cost involved in this strategy – the denial of the transitivity of causation – is no cost at all, since there are no good arguments for the claim that causation is transitive. I also show that another objection of Dowe’s – that the strategy makes causal facts depend on facts that are extrinsic to the causal process in question – fails to hit home.

2. Dowe’s chance-relativising strategy

In this section I argue that Dowe’s strategy for rescuing (IC) does not succeed. I begin by showing that his analysis of a hindering process – the process from which we need to abstract away in order to restore chance increase between *c* and *e* – fails. However, the failure of that analysis does not entail that we should abandon the general chance relativising strategy all together, and I therefore provide a more intuitive, and more successful, conception of the alternative process from which we need to abstract away. I then show that given this conception of the alternative process, the chance relativising strategy can be applied to chance decreasing causal processes that are *not* intuitively cases of causation, for example the defoliant case described above. In other words, all cases of chance decreasing causal processes can be characterised as mixed path cases, and hence all such cases are, according to Dowe’s strategy, cases of causation. So Dowe’s strategy is *too* successful: it makes not just some but all chance decreaseers come out as causes. Hence the general strategy does not succeed because it fails to discriminate between (alleged) chance-lowering causes and chance-lowering non-causes.

At first sight, Dowe’s strategy seems intuitively plausible in the bus-brick case. There, two identifiable and reasonably independent causal processes are going on: the process that involves the bus speeding along the street, and the process that involves the brick falling. My push in effect stops Edna interacting with the first causal process, but at the same time forces her to interact with the second. So – at first sight – it seems plausible to say that the push does two things: it promotes Edna’s death by involving her in the brick process, and hinders her death by getting her out of the way of the bus process.

It also seems to be a straightforward matter to imagine the situation with one of the processes removed: it's easy to imagine the situation minus the bus, where I push Edna (for reasons unknown, or perhaps for no reason), the brick falls, and Edna dies; and it's also easy to imagine the situation minus the brick – where I push her out of the path of the bus and safely onto the (falling-brick-free) pavement. Clearly in the former case, where the bus isn't on the scene, the push increases the chance of Edna's death. Hence it seems plausible to say that the push initiates a 'mixed path' to Edna's death, and that we can sensibly relativise the chance of Edna's death to just one path: the brick-process.

However, we need to be a bit clearer about the nature of the 'paths'; in particular, we need to be precise about the nature of the bus-avoiding 'process' from which we are supposed to be abstracting away. I shall argue that the details of Dowe's analysis do not yield the intuitive picture presented above, and hence that if we want to try to save some form of the chance relativising strategy, we are going to have to hang on to the intuitive picture rather than the details of the analysis.

Picture the scene. Before the push, the bus is hurtling down the street and heading straight for Edna. This is a *bona fide* causal process. Then we have the push. (The bus continues to hurtle down the street, but these later stages of the process are no longer relevant to Edna's death.) The push *stops* Edna from interacting with the genuine causal process of the bus's travelling down the street. However, according to Dowe, when we abstract away from the bus-avoiding 'process', we are supposed to be abstracting away from the hindering 'process' *initiated by the push*. We are not – or at least not directly – supposed to be abstracting away from (the earlier stages of) the genuine causal process of the bus hurtling down the street, since *that* process is neither a potential preventer (i.e. hinderer) of Edna's death, nor a process initiated by the push.

So in what sense does the push initiate a hindering process – a process that could (but in fact does not) lead to Edna's survival? Well, the push prevents a particular course of events – say, bus-being-a-foot-away-from-Edna (event *b*), bus-hitting-Edna (*d*), ..., Edna's death (*e*) – from occurring: without the push, that sequence of events would have been very likely to occur. On Dowe's account, the 'process' *actually* initiated by the push is, as it were, the negation of the earlier stages of that merely possible process $b - d - e$. Call the bus's hurtling down the street prior to the push *a*, and the push *c*. Without *c*, the process that would have been very likely to occur would have been $a - b - d - e$. *With* the push, the bus-avoiding 'process' that *actually* occurs is $a - c - \neg b - \neg d - e$. (Recall that on Dowe's view hindrance is unsuccessful prevention. Had the brick not been there, the hindering 'process' initiated by *c* would have succeeded – it would have prevented Edna's death. In other words, it would have been $c - \neg b - \neg d - \neg e$. But the hindering 'process' was not successful (since Edna in fact dies), hence $c - \neg b - \neg d - e$.) Note that this hindering 'process' is not a *causal* process according to the conception of causal processes presupposed throughout this paper; nor is it a causal process according to Dowe.

Now, Dowe's strategy for saving (IC) is to relativise the chance of *e* to the brick process (call this process ρ). To do this, we need to 'go to the closest ρ -only world, that is to say, the closest world where ρ is the only process between *c* and *e*' (p. 80). In other words, we need to evaluate the chance of *e* at the closest world where the other bus-avoiding 'process' (call it the σ -process) does not occur, but the ρ -process does.

What is such a world like? Well, according to the intuitive characterisation of the situation given earlier, we can think of the closest ρ -only world as one where the bus is simply not on the scene: there, there is (intuitively) no hindering initiated by the push, because the push does not save Edna from any potentially life-threatening road accident, but there is still the genuine causal process from the push, via the brick, to Edna's death. However, this intuitive picture is not what is entailed by Dowe's own analysis. For in such a world, b and d do not happen. So – like the actual world – that world is a world where c , $\neg b$, $\neg d$ and e all 'occur'. But that 'process' is precisely the process – the σ -process – from which we are supposed to be abstracting away.⁹ So the world in which we intuitively *want* to evaluate the chance of e in order to restore chance increase between c and e is *not* a world where the σ -process does not occur. Indeed, the only way of getting to a world where $\neg b$ and $\neg d$ do not occur is to go to a world where b and d *do* occur – that is, a world where Edna *does* get hit by the bus. In such a world, pushing Edna fails to get her out of the path of the bus, and hence, so far as I can tell, makes no difference to her chance of dying.

One reason why Dowe's account fails, then, is that the process we need to abstract away from is not the 'process' – the σ -process – he *says* we need to abstract away from. We really need to go to a world where there is no possibility of Edna getting hit by the bus, for it is only in such a world that the push will fail to be a potential preventer of her death. There are three different kinds of world that satisfy this requirement. The first is a world where early stages of the causal process involving the bus – stages that occur before the push occurs – do not happen, for example a world where there is no bus at all. The second is a world where earlier stages of the bus-process *are* present, but there are other features of the situation that make it impossible for the process to run to completion. One such world would be a world where there is a sufficiently sturdy barrier – a reinforced concrete wall, say – between Edna and the bus. The third and final kind of world is a world where the laws of nature are such that the early stages of the bus-process cannot lead to Edna's death (for example, a world where buses and people repel each other, so that Edna and the bus cannot make contact with each other).

Staying with Dowe's general strategy, we can say that when relativising the chance of e to the brick-process we go to the closest world (of whichever kind) in which the bus cannot knock Edna over. For the purposes of the bus-brick case, it does not matter which kind of world (or which world of a particular kind) we go to, since in all three kinds of world the push does not hinder Edna's death, and hence – relative to the brick-process – c increases the chance of e . (Note, however, that in none of the possible worlds described above does Dowe's bus-avoiding 'process', involving negative events $\neg b$ and $\neg d$, fail to occur – since b and d do not occur in any of them.)

The thought that the bus-brick case is in *some* sense a case of 'mixed paths' thus still seems plausible. So perhaps we could hang on to the chance-relativising strategy in general without adopting Dowe's analysis of the 'process' from which we need to

⁹ Of course, in some intuitive sense the σ -process in the actual world and the σ -process in the bus-free world look very different from one another, since b and d fail to happen in virtue of very different kinds of positive events. Still, assuming negative events occur by definition just if the positive ones don't, they really are the same 'events' at each world, and hence the same 'process'.

abstract away. The basic idea would go something like this: In cases where c lowers the chance of e but nonetheless causes e , e would have some (higher) chance of occurring had c not occurred. So there must have been some way for e to come about without the help of c : there must be some possible causal process – one that does not run to completion in the actual world because of the interference of c – which, if it ran to completion, would cause e . In the bus-brick case, this process is the process involving the bus speeding along the road, hitting Edna, and so on. It is the presence of early stages of this process (plus surrounding circumstances and the laws of nature) in the actual world that makes c lower rather than raise e 's chance, since, were there no such early stages of such a process, or were there a concrete wall in the way, or were the laws different, e would have no way of occurring in the absence of c , and c would automatically raise e 's chance (from zero to something bigger).

In Dowe's terminology, it's the presence of the earlier stages of this potential causal process (plus surrounding circumstances and the laws) that make c *hinder* e : by interrupting the more reliable way of producing e , c 's occurrence has the potential to prevent (that is, successfully hinder) e . So, if we want to restore chance increase between c and e , we need to think of a situation where that alternative potential causal process either does not get off the ground at all, or cannot run to completion. In such a situation, c will *not* be a hinderer of e – not because c somehow fails to initiate a hindering 'process' of negative events, but simply because c no longer has the capacity to prevent e because, without c , e would not occur.

Unfortunately, however, additional problems for this strategy arise when we try to accommodate chance-lowering causes that are not obviously analogous to the bus-brick case. Consider two other cases of chance-lowering causation to which Dowe applies the chance-relativising strategy: the pulled drive case described earlier, and the case of the decaying atom. In the pulled drive case, Sue pulls her drive (c), thereby lowering the chance of the ball landing in the hole (e): had Sue not pulled the drive but instead struck the ball as she had intended, the chance of e would have been higher. In the decaying atom case, an atom can decay to state k (call this event e) via either of two paths, one involving the intermediate product i and one involving intermediate product j . Either way – whether the process runs via i or via j – there is, at the time that i or j occurs, some chance that e will occur; but j gives e a higher chance than i does. Also suppose that time is discrete, and that the relevant steps are right next to each other: there are no relevant events in between the above mentioned steps in the process, nor indeed any times for any such events to occur at. In fact the atom decays via i (let c be the event of its decaying to i); but prior to doing so it could have decayed via j instead. c therefore decreases the chance of e but, according to Dowe, c causes e .

While both of the above cases are similar to the bus-brick case in that c is (allegedly) a cause of e yet lowers e 's chance, they are disanalogous to it in that there is only one genuine causal process going on, rather than two. In each case, there is (intuitively) a single causal process going on, of which c is a part; and c modifies that process in such a way as to lower e 's chance. In the bus-brick case, by contrast, there were two separate processes – one involving the bus and one involving the brick.

How are we to think of the 'mixed paths' in the above cases? Well, in each case (as with the bus-brick case) c not only initiates a genuine causal process (the atom decaying to

state k ; the pulled drive, through the trajectory of the ball to the hole-in-one), but also acts as a hinderer (in Dowe's sense) of e , since c has the capacity to prevent e : had the pulled drive resulted in what it was most likely to result in, it would have prevented a hole-in-one, for example.

Above I argued that the correct way to restore chance increase in the bus-brick case was to go *either* to the closest world where the earlier stages of the process that could, in c 's absence, have led to e do not occur, *or* to the closest world where those earlier stages occur but are somehow incapable of leading to e . So the crucial question is whether we can abstract away from that alternative process in the pulled drive and decaying atom cases too. I shall argue that *if* it is possible to do this in these cases, it is possible to do it in *all* cases of chance decreaseers that are linked by a causal process to their effects. This yields the result that *all* such chance decreaseers are causes – for example, the defoliant causes the plant's survival on Dowe's analysis.

Now, we saw earlier with regard to the bus-brick case that there were two possible ways of abstracting from the alternative process – that is, of going to a world where c is not a potential preventer of e . First, we could go to a world where the actual, earlier stages of the alternative process do not occur (a world where there is no bus at all, for example). Or, second, we could go to a world where those earlier stages *do* occur, but, for some reason, do not have the capacity to cause e (a world with different laws of nature that somehow do not permit the speeding bus to hit Edna, or a world where there is a wall between Edna and the bus, for example).

However, we do not have such a choice in the decaying atom and pulled drive cases: we have to use the second option. This is because earlier, actual stages of the alternative process that has the capacity to cause e are *also* earlier stages of the process which in fact, via c , leads to e . In the pulled drive case, earlier actual stages of the alternative process – the process that could, without c , have led to e – include, for example, Sue putting the golf ball on the tee, lining up for the drive, and taking a swing. But those events are also part of the actual causal process that in fact led, via c , to e . If we go to a world where *those* events do not occur, we'll find that c does not occur either: one cannot pull one's drive without there being a ball to drive or a drive to pull. Similarly for the decaying atom case. The alternative process that might have led to e had c not occurred has as its earlier stages the continued existence of atom h just prior to the time when c occurred; c prevents that process running, via j , to completion. If we go to a world where those actual earlier stages are not present, we go to a world where there is no atom at all, and hence to a world where c cannot occur. Hence worlds where (actual) earlier stages of the alternative process do not occur will not be worlds where c increases the chance of e , since they will be worlds where c does not occur at all.

The reason why we cannot avail ourselves of the first option in these two cases is, of course, the fact that – unlike the bus-brick case – there is only one genuine causal process going on. The earlier stages of the causal process that in fact leads to e consists of the very same events as the earlier stages of the causal process that might, in c 's absence, have led to e . We therefore need to use the second option and go to a world where those earlier stages occur but do not, for some reason, constitute the earlier stages of a process that could, in c 's absence, cause e .

How are we to do this? Well, in the decaying atom case, we have to go to a world where there is only one possible decay path to k – namely the path via i . At such a world – as Dowe points out (p.80) – c really does raise the chance of e , since, at that world, were c not to occur, e would, by stipulation, have no chance of occurring. In the pulled drive case, we have to go to a world where, for some reason, Sue simply has no way of getting a hole-in-one if she doesn't pull her drive. (Imagine, for example, that there is an obstacle that blocks the ball if it's on a straight-drive trajectory to the hole, but not if it's on a pulled-drive trajectory. Or imagine a world where the laws of nature are such that a 'straight' drive would produce a trajectory that would take the ball nowhere near the hole – or a conveniently located tree.)

This sort of world is, I think, the kind of world Dowe has in mind when he talks about going to a world where the hindering process does not occur – though, as I argued earlier, it is not the sort of world where, on his *analysis* of hindering 'processes', the hindering process does not occur. And this strategy really does, as Dowe claims, save (IC) and therewith (CCT), since the worlds described above are indeed worlds where c increases the chance of e . So the chance-relativising strategy really does restore chance increase in cases of causation where c lowers the all-things-considered chance of e .

It seems, then, that although Dowe's own analysis of mixed path cases (according to which the hindering 'process' is construed as a non-causal process involving negative events) does not succeed, his general strategy of abstracting away from the alternative causal process – the one in virtue of which c counts as a hinderer of e – successfully restores (relativised) chance increase in cases of (all-things-considered) chance decreasing causation. However, a successful defence of (IC) that gives credence to common-sense intuitions about chance decreasing causes must yield the result that some, *but not all*, chance decreaseers that are linked to their effects by a genuine causal process are causes of those effects. Recall the defoliant example mentioned earlier: a plant is sprayed with defoliant (c) yet, six months later (after a period of being rather sickly), is in perfect health again (e). c lowers the chance of e , and there is a causal process between c and e . Philosophers who have discussed the case all agree (in a rare case of consensus over a thought experiment concerning causation) that c did not cause e .

Unfortunately, the strategy of abstracting away from the alternative causal process in order to restore chance increase can be applied just as easily to such cases of chance-decreasing causal processes as it can to cases of alleged chance-decreasing causation. For example, the defoliant case is susceptible to just the sort of move made above with respect to the decaying atom and pulled drive cases. The plant would have had a (higher) chance of survival had it not been sprayed (c) because its normally-functioning processes would have been very likely, in the absence of c , to continue to operate and to lead to its survival in the usual way. But we can abstract away from that process, just as we did in the other cases. That is, we can go to a world where, for some reason, failure to be sprayed would lead to the death rather than the survival of the plant. For instance, we can imagine a world where, without the spraying, the plant would succumb to a disease that can only be transmitted through its leaves. At that world, c increases the chance of e , since without c the plant would not survive. So, relativised to the actual causal process between the spraying and the survival, c increases the chance of e – so on Dowe's account, the spraying caused the survival.

It is clear that the chance relativising strategy will generalise to all cases where there is a causal process between c and e . In all cases of causal process plus chance decrease, there must be some alternative potential route that c prevents or cuts short, which might, in the absence of c , have led to e . If there were no such alternative potential route to e , c could not lower the chance of e in the first place, since without c there would be no chance of e occurring at all. Once we see this, and see that the chance relativising strategy in effect abstracts away from the possibility of that alternative causal process running to completion, it is easy, for every chance decreasing causal process, to cook up a possible world where that alternative process is unable to run to completion.

Dowe's attempt to rescue (IC) therefore fails, since it has the consequence that the existence of a causal process between c and e is a sufficient condition for causation: the attempt renders not just some, but *all* cases of chance decreasing causal processes as cases of causation.

3. *Causing, hindering and the bullet-biting strategy*

In my (1997), I present an analysis of causation according to which hindrance is a kind of causal relation. Ignoring the question of how a causal process is to be defined – an issue that is not directly relevant to the purposes of this paper – the analysis's central principle is the following:

- (H) c and e are causally related if and only if there is a causal process between them. If there is a causal process between c and e , then c causes e if and only if c increases the chance of e , and c hinders e if and only if c decreases the chance of e .¹⁰

(H) entails that there is no such thing as chance-decreasing causation, since it entails that alleged cases of chance-decreasing causation, like the pulled drive case, the bus-brick case and the decaying atom case, are in fact cases of hindrance rather than causation. My strategy for rescuing (IC) is therefore the rather simplistic strategy of biting the bullet.

In this section I argue that, if we accept (H), the bullet-biting strategy is a perfectly reasonable strategy, since (H) provides the resources to take the pain out of the bite. First, however, I show that analyses of indeterministic causation that do not take hindrance to be a species of causal relation face a serious problem that does not arise for (H), and that such analyses do justice to intuitions about alleged chance-lowering causation at the expense of violating intuitions about chance-lowering causal processes that are *not* cases of causation – the defoliant case, for example. I also offer a somewhat speculative diagnosis of why standard analyses of causation fail to accord hindrance the status of a causal relation and argue that the reasons for failing to do so are bad reasons. In the next section, section 4, I argue that, once (H) is accepted, some objections raised by Dowe to the bullet-biting strategy can be answered.

¹⁰ As stated, the analysis says nothing about the case where there is a causal process between c and e but c does not change the chance of e at all, and thus remains silent for cases of deterministic pre-emption where c leaves the chance of e at 1 – just what it would have been in the absence of c . This omission can be rectified by stipulating that c causes e if and only if there is a causal process between c and e and c does not decrease the chance of e .

I argued in section 2 that Dowe's attempt to restore chance increase to cases of chance-decreasing causation fails because it makes *all* chance-decreasing causal processes come out as cases of causation. In fact, it is not very surprising that Dowe's analysis should have as a consequence that all causal processes – whether chance increasing or chance decreasing – turn out to be cases of causation, since for Dowe hindering is not a species of causal relation: hindering 'processes' are not *causal* processes. Roughly speaking, for Dowe *c* hinders *e* if and only if *c* cuts off some process which would, had *c* not occurred, have been more likely to cause *e*. So *c* hinders *e* not in virtue of the *actual* causal process that runs from *c* to *e*, but in virtue of the fact that *c* stops some *other* causal process from running to completion. Since there is, by stipulation, a causal process between *c* and *e*, and since the fact that *c* hinders *e* obtains not in virtue of the existence of *that* causal process but rather in virtue of the non-existence of some other causal process, when we ask what the *causal* relation is between *c* and *e* – the relation that obtains in virtue of the existence of a causal process between *c* and *e* – the only available answer is that *c* causes *e*.

Other analyses similarly have no room for hindrance as a species of causal relation. Lewis (1986), for example, defines causal dependence as chance increase, and then defines causation as a chain of causal dependence. This has the effect of 'factoring out' any chance decrease: when *c* lowers the chance of *e* there is generally some chain of events (*d* and *f*, say) such that *c* increases the chance of *d*, *d* increases the chance of *f*, and *f* increases the chance of *e*. In all such cases – the defoliant case, for example – *c* comes out as a cause of *e*. Menzies' (1989) analysis produces the same result. Dowe calls such accounts 'interpolating' accounts because they attempt to hook a chance-decreasing cause to its effect by interpolating further, chance-increasing events between cause and effect.

Interpolating accounts in effect render all causal processes as cases of causation; like Dowe's 'mixed paths' analysis, they turn the causal process condition into a sufficient condition for causation. Hence they only yield the 'right' result in the pulled drive and bus-brick cases at the expense of making all chance decreasing causal processes come out as cases of causation. So they save intuitions about chance-decreasing causes at the expense of violating intuitions about chance-decreasing non-causes – for example the intuition that the spraying was not a cause of the plant's survival.

Moreover, interpolating accounts run into worse trouble in alleged cases of 'direct' causal relations. In Dowe's decaying atom case, for instance, we are supposed to imagine that time is discrete, so that no further events can be interpolated between *c* and *e* in such a way as to yield a chain of chance increasing causal dependence. Since interpolating accounts recognise only one kind of causal relation – causation – they do not have the resources to recognise any kind of causal relation at all between *c* and *e* in the decaying atom case, rendering *c* and *e* completely causally unrelated. Recognising hindrance as a species of causal relation, however, provides a simple solution to the problem posed by the decaying atom case, since (given an appropriate characterisation of a causal process) (H) allows us to say that *c* and *e* are causally related, since *c* hinders *e*.¹¹

¹¹ I present the argument that interpolating accounts cannot deal with direct chance-decreasing causal relations in more detail in my 1997.

Why is it that analyses of causation typically recognise only one kind of causal relation – namely causation? Well, one way of explaining it is to see it as a hangover from deterministic analyses of causation. Under the assumption of determinism, there really is no hindrance (construed as a chance-decreasing kind of causal relation), since to lower e 's chance under determinism is to lower it from 1 to 0; hence it is impossible for c to decrease the chance of e and for e still to occur. Once we abandon determinism, however, there is no good reason for thinking that genuine hindrance does not exist, or that it is somehow not a genuinely causal relation. With indeterminism in place, there is no reason to regard chance decrease as any less real or any less important than chance increase when it comes to analysing causation, and hence no reason to regard the kind of causal relation manifested by chance-decreasing causal processes – hindrance – as any less real or important than the kind – causation – manifested by chance-increasing processes.

One might object that the term 'hinders' can perfectly appropriately be used in deterministic contexts, and hence that hindrance cannot (as I've claimed) be a feature of the world that only exists if indeterminism is true. For example, suppose I succeed in lighting a damp match. Whether or not the relevant processes are deterministic, it seems appropriate to say that the match's dampness hindered its lighting. But on my account of hindrance, in the deterministic case the dampness didn't *really* hinder the lighting. When I struck the match, circumstances were such as to guarantee that the match would light. The dampness of the match may or may not have been a necessary condition of the match's lighting. If it was, then the dampness was in fact a cause of the lighting, and if it wasn't, then the dampness was simply irrelevant to the lighting. Hence (so the objection goes) whatever the relation is that I've called 'hindrance', it cannot really *be* hindrance, since intuition tells us that the term 'hindered' applies in cases – namely, deterministic cases – where on my account the term does not apply.

This objection is not a serious one, since so far as I can tell, the desire to say that the dampness hindered the lighting simply goes away if one takes the deterministic starting point of the objection seriously. If the dampness was necessary for the lighting (say because I only struck with sufficient force because I knew the match was damp – if it hadn't been damp, I would not have taken such care, and it would not have lit), then it seems perfectly appropriate to say that the dampness caused, rather than hindered, the lighting. And if the dampness made no difference to whether or not the match lit, it seems perfectly appropriate to say that the dampness was irrelevant to the lighting.

In fact, my use of the term 'hindered' does not differ so very much from the deterministic usage described above. According to (H), hindrance is an indeterministic relation characterised by chance decrease. Deterministic usage of the expression 'hindered' can be seen as expressing the idea that c lowers the *probability* of e , where the probability of e is a probability of the non-single-case variety – the kind that can take values other than 0 and 1 even if determinism is true. Generally speaking, damp matches light far less frequently than dry ones; hence, given some incomplete specification of the circumstances in which I strike the match, the dampness lowers the probability (though not the chance) that the match will light. Of course, the non-single-case probability of e will vary depending on how the situation is described, and hence there is no univocal answer to the question 'did c hinder e ?' in deterministic situations. This is not the case

with hindrance as defined by (H), since single-case chances are not relative to how the relevant events are described.

We are still left with the problem set up at the beginning of the paper – that of chance-decreasing causation – since (H) is incompatible with its existence. According to (H), Sue’s pulled drive hindered (and did not cause) the hole-in-one, the atom’s decaying to state i hindered (and did not cause) its decaying to state k ; and the push hindered (and did not cause) Edna’s untimely death.

According to Dowe (and no doubt according to others too), such results go against our intuitions: intuitively, in all three cases, c caused e . I do not think the case for respecting intuitions is particularly strong here. Recall that on standard analyses of causation, there is no distinction between ‘ c and e are causally related’ (or ‘there is a causal process between c and e ’) and ‘ c causes e ’. So, on such analyses, to deny that c caused e is to deny that there is any kind of causal relation or connection whatsoever between c and e . Given this starting point, the desire to say that c caused e in the above cases is natural, since in those cases it would indeed be implausible to claim that c bears no causal relation to e at all. However, once we hold that causation is not the only kind of causal relation, we can safely deny that c caused e without rendering c and e causally unrelated. And this is precisely what taking hindrance to be a species of causal relation allows us to do: to say that c hindered e is to assert that c and e are causally related; but it is also to deny that the causal relation thereby instantiated is the relation of causation.

One might nonetheless claim that brute intuitions in the three cases of alleged chance lowering causation should be respected. Perhaps if there were some viable analysis of causation that yielded the result that c caused e in all three cases, that would be a good reason to prefer that analysis. However, as I have argued, the prospects for such an analysis are dim.

4. Dowe’s objections to the bullet-biting strategy

Dowe raises two objections against the bullet-biting strategy of denying that alleged chance-decreasing causes really are causes. First, he points out that the bullet-biting strategy entails that causation is not transitive, since according to that strategy it can be the case that c causes d , d causes e , but c hinders (rather than causes) e . The bus-brick case is an example of this: the push causes Edna to be on the pavement, which in turn causes her to be hit by the brick, which in turn causes her death. But the push hinders, rather than causes, the death.

It is certainly true that most philosophers are very keen to hang on to the transitivity of causation. Douglas Ehring, for example, claims that ‘transitivity is a fundamental logical feature of the causal relation ... causal transitivity should be disavowed only as a last resort’¹². However, I have not been able to find any *arguments* for the thesis.¹³ Perhaps

¹² Ehring (1997), p. 82

¹³ In a recent paper, E.J. Hall (2000) also claims that transitivity is too important to give up – and that we should only do so if we have independent reason for it. However, Hall is concerned solely with deterministic causation, where there is no distinction between the existence of a causal process and

one reason why transitivity is so popular is that it functions as a kind of methodological principle. When we want to trace the causes of some event – Edna’s death (*e*), say – we typically start by identifying the event’s immediate causes (being hit by the brick (*d*), say), and then tracing back to the causes of *those* events, and so on. It seems that if causation were not transitive, there would be no guarantee that such a procedure would reveal distant causes of Edna’s death, since without transitivity the fact that each step was caused by the preceding step is no guarantee that the first step caused the last.

Does denying transitivity amount to denying the general applicability of this valuable methodological principle? Well, no – not on the view proposed here. On my view, the ‘there is a causal process between’ relation *is* transitive. So, in tracing the steps in the chain of events that led to Edna’s death, we are identifying the *causal process* that led to it. For example, when we establish that the push (*c*) caused *d*, and that *d* caused *e*, we have – by the transitivity of causal processes – established that there is a causal process between *c* and *e*. The only difference is that, before jumping to the conclusion that *c* caused *e*, we have to determine whether or not *c* raised the chance of *e*. In the current case, it turns out that *c* did not raise *e*’s chance; hence we should conclude that *c* hindered *e*.

It’s worth reiterating the point that hindrance (so defined) is a phenomenon that can only occur if the world is indeterministic. In deterministic situations, the existence of a causal process between *c* and *e* guarantees that *c* causes *e*, since it is impossible for there to be a causal process between *c* and *e* if *c* lowers the chance of *e* from 1 to 0. Given that many analyses of causation presuppose determinism, or are descendants of analyses that presuppose determinism, it is therefore not surprising that it should be so widely taken for granted that causation is transitive – since, assuming determinism, there is a causal process between *c* and *e* if and only if *c* causes *e*.¹⁴ It is only when we abandon determinism that causation and causal processes come apart and we can legitimately ask whether one of the two relations might not be transitive. And I can see no reason not to deny the transitivity of causation so long as we hold on to the transitivity of causal processes. Indeed, giving up on the transitivity of causation is the only way of getting the right result in the defoliant case. There is clearly a causal process running from the spraying to the survival – which is to say, there is a chain of events running from the spraying to the survival such that each event in the chain caused the next. So the only way to deny that the spraying caused the survival is to deny that the existence of that chain of causation is sufficient for the first step to cause the last.

Dowe’s other objection is that if we deny that alleged cases of chance decreasing causation really *are* cases of causation, we make the obtaining of the causal relation depend on extrinsic features of the situation; and, intuitively, causation is an intrinsic matter. As Peter Menzies puts it:

I drop a piece of sodium into a beaker of acid and that causes an explosion, this causal relation is an intrinsic feature of the cause-effect pair. So if there is another person waiting in the wings, ready to drop a piece of sodium into the beaker if I do not, that

causation. I think the move to indeterminism *does* provide an independent reason to give up the transitivity of causation – but not, as we shall see, the transitivity of causal processes.

¹⁴ Recall that I am presupposing that (CP) is true.

makes no difference to whether the causal relation holds between my dropping the sodium and the explosion. The presence of the alternative cause is neither here nor there to the causal relation that exists between the actual cause and effect. The causal relation does not depend on any other events occurring in the neighbourhood: the causal relation is intrinsic, in some sense, to its relata and the process connecting them.¹⁵

Dowe's objection is that denying that, for example, the bus-brick case is a case of causation amounts to denying that the causal relation does not depend on any other events occurring in the neighbourhood, and therefore makes causation unacceptably extrinsic. He says:

... it seems implausible to suppose that whether the push caused the death is dependent on how fast a bus was going, a bus what's more that didn't hit her. Suppose the chance that the lady dies given the push is 0.5. Suppose if the bus is travelling at 36 miles per hour that the chance that the old lady will die given that there is no push is 0.45, and that if the bus is travelling at 38 miles per hour that the chance that the old lady will die given that there is no push is 0.55. Then, if the bus is travelling at 36 miles per hour when the push occurs, which leads to her being hit on the head by a brick and dying, then the push caused her death. But if the bus is travelling at 38 miles per hour when the push occurs, which leads to her being hit on the head by a brick and dying, then the push does not cause her death. The intuition is that whether an event causes another via a particular process shouldn't depend on the strength of a separate, distant unsuccessful process.¹⁶

The issue of whether causation is an intrinsic relation is a difficult and controversial one (not least because one's answer depends on how one defines 'intrinsic'¹⁷). I shall, however, ignore the technicalities of this debate and assume some sort of pre-theoretic understanding of 'intrinsic' and 'extrinsic', according to which, for example, how fast the bus is going is clearly extrinsic to the causal process initiated by the push, whereas features like the strength of the push and the velocity of the brick are intrinsic features of the process. I shall argue that both Menzies and Dowe are here running together what are, on my view, two separate questions: the question of whether the *causal process between* relation is intrinsic, and the question of whether the *causal* relation is intrinsic.

Before doing that, however, let's look at the broader picture for a moment. There is a general tension between, on the one hand, thinking of causation as an intrinsic relation and, on the other, holding (as most contemporary theorists do) that an adequate analysis of causation must take account of the chances that causes give their effects. (This is meant to include not just theories that appeal directly to chance in their analysis of causation, but also to theories that simply seek to do justice to the thought that causation and chance are somehow conceptually connected. For example, a commitment to (CCT) displays a commitment to a connection between causation and chance, whether or not (CCT) actually plays a part in one's analysis of causation.)

This tension arises because whether or not *c* raises the chance of *e* depends not just on the intrinsic features of *c* – and not even just on the intrinsic features of *c*, plus the intrinsic features of the causal process running from *c* to *e*, plus the laws of nature – but on facts

¹⁵ Menzies (1998), p. 339

¹⁶ Dowe (2000) p. 75

¹⁷ See for example Langton and Lewis (1998).

that are entirely extrinsic to c and to the causal process that thereby leads to e . In the bus-brick case, for example – as Dowe notes – c lowers rather than raises the chance of e because of the speed at which the bus happens to be going. In the defoliant case, c lowers rather than raises the chance of e because in fact external factors are not such as to jeopardise seriously the plant's chances of survival were all its leaves to remain intact. Whether or not c raises the chance of e depends in part on what the chance of e would have been in the absence of c ; and that generally depends on features of the world that are nothing to do with the intrinsic features of the actual causal process between c and e .

Moreover, whether c raises or lowers the chance of e also depends on what the *actual* chance of e is; and this also often depends on facts that are extrinsic to the causal process linking c and e . Suppose, for example, that the defoliant only has a 50% chance of affecting the plant's leaves at all. Then the plant's *actual* chance of survival just after it is sprayed depends not only on how likely it is that the plant will die if it loses its leaves, but also on how likely it is that it will die if its leaves remain intact. Suppose that circumstances are such that if the plant keeps its leaves (whether or not it is sprayed) it will have a 90% chance of survival, but that if it loses its leaves (which it can only do if sprayed) its chance of survival will be just 20%. Then just after spraying, the plant's actual chance of survival is 55%. Now suppose that circumstances are such that if it keeps its leaves it will have only a 10% chance of survival (because of the presence in its immediate environment of a nasty leaf disease, say). Then just after spraying, the plant's actual chance of survival is 15%. Thus extrinsic facts – facts about the surrounding environment – help determine the actual chance of e , and thus whether or not c raises the chance of e , even though those extrinsic facts play no part in the actual causal process leading from the spraying to the survival. It is therefore very difficult to hold on to the idea that causation is an intrinsic matter while at the same time holding that causation and chance are related.

On the view of causation as a chance-increasing causal process, causation is (at least partially) extrinsic in a very obvious way, since, as we saw above, whether or not c increases the chance of e is at least partially an extrinsic matter. Is this a counter-intuitive result? I think not. For one thing, intuitions about whether c causes e sometimes *do* hinge on extrinsic features of the situation. So to rule that causation must *never* depend on extrinsic features would violate those intuitions. Consider the defoliant case. As the case is actually set up (without the story about the leaf-infecting disease), common-sense intuition favours the verdict that the spraying was not a cause of the plant's survival. But now consider the variant case where, in the absence of the spraying, the plant would be highly susceptible to a fatal disease. In *that* case, I think common-sense intuition favours the verdict that the spraying *was* a cause of the plant's survival. Yet the differences between the original case and the variant case are purely extrinsic differences: the intrinsic nature of the actual causal process leading from the spraying to the survival are precisely the same in each case.

Second, given (H), we must remember that on the account offered here, the question of whether c causes e must be distinguished from the question of whether there is a causal process running from c to e . With the distinction in place, we can accept that causation itself is (at least partially) extrinsic without being required to accept that the existence of a causal process between c and e is likewise an extrinsic matter. That is, the extrinsicity

of causation is perfectly compatible with a fully intrinsic characterisation of the notion of a causal process.¹⁸ Granted, on my view whether or not the push is a cause of Edna's death depends in part on how fast the bus is going. But whether there is a causal *process* between the push and the death does not depend on how fast the bus is going, nor even on whether there is a bus on the scene at all.

Both the transitivity objection and the intrinsicity objection can be met, then, by separating questions about causation from questions about causal processes, and showing that according to (H), the causal-process-between relation can meet the desiderata even though causation itself does not. For the objections to stick, one would have to show that, once the two relations have come apart, it is causation, rather than (or perhaps as well as) the causal-process-between relation to which transitivity and intrinsicity properly apply. I can see no reason to think that this can be done. Once (H) is accepted, then, the bullet-biting strategy is a plausible strategy for saving (IC): we can retain the view that causes increase the chances of effects without appealing to the chance-relativising strategy.

References

- Cartwright, N. (1979), 'Causal Laws and Effective Strategies', *Nous* 13 (1979), pp. 419-437
- Dowe, P. (1992), 'Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory', *Philosophy of Science* 59, pp. 195-216
- Dowe, P. (2000), 'Causing, Promoting, Preventing, Hindering' in Ledwig, M., Spohn, W., and Esfeld, M. (eds.), *Current Issues in Causation*. Paderborn: Mentis-Verlag, pp. 69-84.
- Beebe, H. (1997), 'Taking Hindrance Seriously', *Philosophical Studies* 88, pp. 59-79
- Beebe, H. (forthcoming), 'Causing and Nothingness', in J. Collins, E.J. Hall and L.A. Paul (eds.), *Causes and Counterfactuals*, MIT Press.
- Ehring, D. (1997), *Causation and Persistence*, New York: Oxford University Press
- Hall, E.J. (2000), 'Causation and the Price of Transitivity', *Journal of Philosophy* 97, pp. 198-222
- Langton, R. & D.K. Lewis (1998), 'Defining Intrinsic', *Philosophy and Phenomenological Research* 58, pp. 333-345
- Lewis, D.K. (1986), 'Postscripts to "Causation"' in his *Philosophical Papers, volume II*. New York: Oxford University Press, pp. 172-213
- Mellor, D.H. (1995), *The Facts of Causation*. London: Routledge.
- Menzies, P. (1989), 'Probabilistic Causation and Causal Processes: A Critique of Lewis', *Philosophy of Science* 56, pp. 642-663
- Menzies, P. (1998), 'How Justified are Humean Doubts about Intrinsic Causal Relations?', *Communication and Cognition* 31, pp. 339-364

¹⁸ In fact, on my own view (indeed on most views), causal processes are not wholly intrinsic either. But they are at least *reasonably* intrinsic in the sense that whether there is a causal process between, say, the dropping of the sodium and the explosion does not depend on whether there is someone waiting in the wings to drop the sodium in if I do not; and whether or not there is a causal process between the push and Edna's death does not depend on how fast the bus is going.