

Causes of causes

Alex Broadbent

November 27, 2010

Pre-print (Nov 2010) of a paper forthcoming in Philosophical Studies.

Abstract

When is a cause of a cause of an effect also a cause of that effect? The right answer is either “Sometimes” or “Always”. In favour of “Always”, transitivity is considered by some to be necessary for distinguishing causes from redundant non-causal events. Moreover transitivity may be motivated by an interest in an unselective notion of causation, untroubled by principles of invidious discrimination. And causal relations appear to “add up” like transitive relations, so that the obtaining of the overarching relation is not independent of the obtaining of the intermediaries. On the other hand, in favour of “Sometimes”, often we seem not to treat events that are very spatiotemporally remote from an effect as its causes, even when connected to the effect in question by a chain of counterfactual or chance-raising dependence. Moreover cases of double prevention provide counterexamples to causal transitivity even over short chains. According to the argument of this paper, causation is non-transitive. Transitivity provides no viable account of causal redundancy. An unselective approach to causation may motivate resisting the “distance” counterexamples to transitivity, but it does not help with double prevention, and even makes it more intractable. The strongest point in favour of transitivity is the adding up of causal relations, and this is the point that extant non-transitivity analyses have not adequately addressed. I propose a necessary condition on causation that explains the adding up phenomenon. In doing so it also provides a unifying explanation of distance and double prevention counterexamples to transitivity.

1 Options

When is a cause of a cause of an effect also a cause of that effect? There are three options.

1. Always. Causation is transitive. If c causes d and d causes e , then c causes e . A relation R is transitive if and only if $\forall(x, y, z)((Rxy \& Rxz) \supset Rxz)$. An example of a transitive relation is “being taller than”.
2. Not always. Causation is non-transitive. It is not the case that, if c causes d and d causes e , then c causes e . A relation R is non-transitive if and only if $\neg\forall(x, y, z)((Rxy \& Rxz) \supset Rxz)$. An example of a non-transitive relation is “touching”. It is not the case that, for all x , if x is touching y and y is touching z , then x is touching z .

3. Never. Causation is intransitive. If c causes d and d causes e , then it is not the case that c causes e . A relation R is transitive if and only if $\forall(x, y, z)((Rxy \& Rxz) \supset \neg Rxz)$. Intransitivity is thus a species of non-transitivity (intransitivity entails non-transitivity but not vice versa). Among cubes of the same size, the relation of having a whole face touching is intransitive: if two cubes have a whole face touching, then a third cube cannot have a whole face touching a whole face of both the first two.

I am not aware that anyone has found the third option plausible. There seem to be plenty of counterexamples. I strike the white ball, which hits the red, which rolls into the pocket. Striking the white ball caused it to hit the red, and that caused the red to roll into the pocket; it is also the case that striking the white ball caused the red to roll into the pocket. That is why I played the shot, and why I am so pleased with it.

It seems, then, that there are two serious possible answers to the question when a cause of a cause of an effect is a cause of that effect: always; or only sometimes. In this paper, I shall defend the latter answer: I shall argue that causation is non-transitive (but not intransitive).

Some philosophers have thought that transitivity is a necessary ingredient in any adequate account of causal redundancy. Section 2 argues that this is a poor reason for making causation transitive: the solutions that rely on this ingredient are not sufficiently convincing, and the approach puts the evidential cart before the theoretical horse. A more compelling argument for transitivity is that it is a corollary of an “unselective” conception of causation, which is the conception that most — though not all — philosophers take to be the proper object of their attention. Section 3 distinguishes two cases in which causal transitivity appears to fail: spatiotemporal distance between cause and effect, and double prevention. I go on to argue that adherence to an unselective notion of causation addresses only distance counterexamples. Worse, it appears to make double prevention more intractable. In Section 4 I present what I take to be the strongest argument in favour of causal transitivity, namely the fact that causal relations “add up” like transitive ones. When c causes d and d causes e , and it is also the case that c causes e , that is typically not logically independent of the intermediary causal relations. I suggest that existing approaches making causation non-transitive do not offer any adequate alternative explanation for this logical independence. Section 5 distinguishes two kinds of difference making, *cause-led* and *effect-led*. The latter is familiar in the literature on contrastive explanation but has been widely neglected as a potential analysis for causation itself. Section 6 argues that effect-led difference-making is logically implied in those cases where causation behaves like a transitive relation. If we take effect-led difference-making to be a necessary condition on causation then this seems to explain the additive properties of causal relations without recourse to transitivity. Moreover, by identifying the circumstances under which this necessary condition is implied, it is possible to provide a unifying explanation of distance and double prevention failures of transitivity.

2 Redundacy

One common but unsatisfying reason for supposing that causation is transitive is the view that transitivity is required to frame a plausible necessary condition on causation using counterfactuals. Suppose we begin with this claim:

$$c \text{ causes } e \text{ iff } \neg Oc \Box \rightarrow \neg Oe$$

Assume that c and e are actual and distinct events.¹ “ Oc ” means that c occurs, and “ $\Box \rightarrow$ ” indicates the counterfactual conditional.² On this view, counterfactual dependence is both necessary and sufficient for causation. However, the necessary condition thus imposed is generally thought to be implausible, because it means that our common sense judgments about causation are incorrect in relatively mundane cases of causal redundancy. Suppose Able and Baker are throwing rocks at a bottle. Able throws, and smashes the bottle; had he not thrown, Baker, also an excellent shot, would have thrown. The bottle might still have smashed (we need not assert that it *would* still have smashed — only that it is not the case that it would not, i.e. that it might). In this case, our simple theory tells us to deny what is obviously true, that Able’s throw caused the bottle to smash.

The simple theory does have its defenders (e.g. Coady 2004).³ However a much more common reaction is to seek a replacement necessary condition on causation. Lewis’s initial suggestion is that, for c to cause e , it is necessary that c and e be distinct events, and that e counterfactually depends on c , or on an intermediate event d_1 which depends on c , or on d_2 which depends on d_1 which depends on c , or... (Lewis 1973a: 167).⁴ In short, it is necessary that c be connected to e by a chain of counterfactual dependence, of at least one step.⁵

¹I assume that causation is a relation on events, despite being attracted to the view that it is the truth maker for a connective of some sort (Mellor 1995, 2004). For the purposes of discussing transitivity, the difference does not matter, so long as whatever you think causation is supports a corollary of transitivity in the first place. Mellor’s connective clearly does.

²Lewis’s original discussion is framed in terms of counterfactual dependence, which imposes the further requirement that $O(c) \Box \rightarrow O(e)$. But this follows on Lewis’s semantics for counterfactuals, specifically the Centering Assumption, from the requirement that c and e actually occur (Lewis 1973c: 26–31). The Centering Assumption is not uncontroversial (cf. Nozick 1981; Menzies 2004). Although I am not committed to Lewis’s semantics, I will not challenge the Centering Assumption. Hence it is safe to ignore this point for present purposes, and to assume that — for any pair of events x and y which satisfy the existence and distinctness requirements on causal relata — y counterfactually depends on x iff $\neg O(x) \Box \rightarrow \neg O(y)$. Note that Lewis also ignores the complications of his original treatment in later work (Lewis 2004: 78; cf. Lewis 1973a: 164–167).

³Coady argues that for a more precise individuation of events. If Able had not thrown, then that exact bottle smash would not have occurred; a slightly different one would have occurred in its place. However, this strategy is open to a class of objections more commonly made against Lewis’s influence account — namely, that sometimes *exactly* the same effect may be brought about in more than one way. Hall, crediting Yablo, imagines that Able has got his hands on a Smart Rock, programmed to hit with exactly the same velocity at exactly the same moment, regardless of small differences in throw (Hall 2004b: 237). And I have given the less far-fetched example of an alarm clock, which goes off at 7am regardless of whether I press the button to set it, or my wife does (Broadbent 2007: 177).

⁴His second solution, in terms of quasi-dependence, arguably does not rely on the assumption that causation is transitive, as discussed below. But his third solution in terms of influence does (Lewis 2004).

⁵I eschew Lewis’s terminology of causal dependence because it seems to me obviously prejudicial. Who could bring themselves to deny that two events connected by a chain of causal dependence are related as

Suppose we break down Able's activities into a chain of three events. Able throws, the rock strikes the bottle, and the bottle smashes. The point of involving chains is to claim that, if the throw had not occurred then the impact would not have occurred, and if the impact had not occurred then the smash would not have occurred. The hope is that these counterfactuals remain true even though it is false that, if Able had not thrown, the bottle would not have smashed. It is thus central to the defence of the transitivity of causation that counterfactuals themselves are not transitive (Lewis 1973c: 32–35; Lewis 1973a: 167).

There is an obvious objection: the final step in the chain seems to be false. Granted, if Able had not thrown then the rock would not have hit the bottle; but if the rock had not hit the bottle, then surely, we might object, the bottle would still have smashed, thanks to Baker's throw? The further ingredient which Lewis introduces to prevent this move is a ban on *backtracking counterfactuals*, that is, counterfactuals whose antecedents refer to matters at a time later than the matters referred to in the consequent. Although the objection itself does not mention a backtracker, the reasoning behind it involves backtracking. According to Lewis, we reason that, if the impact had not occurred, that would be because Able would not have thrown (a backtracker); so Baker would have thrown instead, and the bottle would still have smashed (Lewis 1973a: 172). The first step in our reasoning is a backtracking counterfactual, and backtracking reasoning is fallacious (Lewis 1979: 33–35).

So, at least, Lewis would argue, in defence of his necessary condition on causation. But is this a good reason to insist that causation is transitive? No, for three reasons.

First, the solution puts the theoretical cart before the evidential horse, to the extent that it is legitimate to think of our untutored intuitions about causation as evidence to which philosophical theories are accountable. We have the intuition that Able's throw caused the smash and Baker's did not. But we also have the intuition that Able's birth did not cause the smash, even if there is a causal chain from that event to the smash. Sacrificing one intuition for the sake of another is irksome enough: but sacrificing one intuition for the sake of a *theory* about another is usually intolerable, unless the theory in question is extremely powerful.

Second, the solution does not work. A large literature proposes much more difficult cases than the easy one I have described. But even in this easy case, the solution is weak, because it attributes to the objector a particular line of reasoning — a backtracking line. It is, to say the least, questionable whether a semantics for counterfactuals can or ought to deny backtracking in such a thorough-going way (a point we will return to in Section 5). But even setting that thorny issue aside, why must the objector backtrack? All she needs to maintain is that, had Able's rock not hit the bottle, the bottle might still have smashed.⁶ One way to do this is the backtracking route Lewis describes, supposing away Able's throw. But another is to suppose that Able threw and missed. Then, without backtracking, the objector can claim that Baker might have thrown and that the bottle might still have smashed. — Later, to be sure: but so much later as to be a different event? It is hard to believe that the smash that might have oc-

cause to effect? Arguably, using "dependence" to express the truth of a counterfactual is also prejudicial, since "depends" inherits heavy causal overtones from its everyday meaning; but let us put up with that for the moment.

⁶Which is enough to negate the claim that it would not have (Lewis 1973b: 2).

curred had Able thrown and missed is a different event from the actual smash, when we have already conceded that the smash that might have occurred had Able not thrown at all is the same event as the actual smash.

Third, even for those who are wedded to counterfactual accounts of causation, or to difference-making accounts more generally, there are other strategies for handling redundancy which do not require causation to be transitive, despite what is sometimes said to the contrary (cf. Paul 2004: 206). One of these is the “contingency” strategy pursued by Joseph Halpern and Judea Pearl, according to which:

...while effects may not always counterfactually depend on their causes in the actual situation, they do depend on them under certain contingencies. (Halpern and Pearl 2005: 844)

David Lewis’s tried-and-rejected notion of quasi-dependence is related to this idea, the “contingency” being a non-actual intrinsic duplicate of the actual situation without the pre-empted event (Lewis 1986b: 206). Christopher Hitchcock also pursues the contingencies strategy (Hitchcock 2001, 2007 — the latter explicitly relying on an account along the lines of the former or of Halpern and Pearl’s to provide clear intuitions in cases of preemption: 519). Another alternative is the strategy I have pursued elsewhere, of arguing that there are other counterfactuals whose truth — unlike that of the dependence of effect on cause — survives in cases of redundancy (Broadbent 2007). And there may yet be others. Insisting that causation is transitive is not compulsory for counterfactual, or more broadly difference-making, accounts of causal redundancy.

In short: transitizing causation for the sake of an account of causal redundancy is a bad move in principle, yielding a poor account of redundancy in practice, and moreover is unnecessary even within the framework of a counterfactual, or more broadly difference-making, approach to causation.

3 Selection

Causal selection is an umbrella term for our various practices of picking out one event, or a handful, from many that are causally connected to a given effect, and honoring it by referring to it as “the” cause (or them as “the” causes). Often we do a lot more than that: we predict, explain, manipulate, assign moral blame and found legal judgements on this sort of singling out. Many philosophers view this practice as extraneous to causal judgment, something we often do *alongside* judging that *c* causes *e* but which is no part of the *judgment* that *c* causes *e*. On this view, causal selection is a choice made by us, in view of our interests, motives, morals, and goodness knows what else, among events which are causally on a par:

We sometimes single out one among all the causes of some event and call it “the” cause. Or we single out a few as the “causes”, calling the rest mere “causal factors” or “causal conditions.” Or we speak of the “decisive” or “real” or “principal” cause... I have nothing to say about these principles of invidious discrimination. I am concerned with the prior question of what it is to be one of the causes (unselectively speaking). My analysis is meant to capture a broad and non-discriminatory notion of causation.

(Lewis 1973a, 162)

And moreover, this attitude to causal selection can motivate the assertion that causation is transitive. Thus Lewis writes:

We have the icy road, the bald tire, the drunk driver, the blind corner, the approaching car, and more. Together, these cause the crash... But these are by no means all the causes of the crash. ...each of these causes in turn has its causes; and these too are causes of the crash. (Lewis 1986a: 214)

Here, the causes of the causes of the crash are assimilated to the general constellation of events in the causal history of the crash; and, in non-discriminatory spirit, Lewis wants to call all of them causes. Ned Hall expresses the point like this:

...transitivity helps to make for an egalitarian relation: Events causally remote from a given effect will typically not be *salient* — but will still be among its causes for all that. (Hall 2004a: 112)

On this view, transitivity is a natural extension of the widely accepted doctrine that causation ought to be sharply distinguished from causal *selection*.

The exact nature of the link between unselective causation and transitive causation is hard to pin down. If you accept that counterfactual dependence is sufficient for causation then you will sometimes need to accept that events very spatiotemporally distant to a given effect are among its causes — my birth, or the Big Bang, as causing me to write these words, for example — and that in turn might lead you to doubt that spatiotemporal distance provides any good objective basis for denying that the causal relation obtains. (Otherwise one would seem committed to the implausible view that, in principle, the existence of a back-up event c_2 which will bring about e absent c_1 can make a difference to whether c_1 causes e event if c_2 is causally independent of both c_1 and e .)

This kind of worry is not confined to counterfactual analyses of causation. One might wonder quite generally exactly where the line is drawn between events that are proximate enough to count as causes, and event that are too distant. If the first event in a chain of events causes the last, then what objective basis is there to say that the first event also causes the effects of the last event? This might lead one to suppose that the reason we usually do not mention very distant events as causes is the same as the reason that we do not mention many other causally relevant factors as causes: because we don't consider them salient. But, we might suppose, they are causes nonetheless.

As a motivation for declaring causation transitive, this line of thought suffers from three defects. First, most obviously, the link between selection and transitivity remains vague. It is *possible* to retain an unselective notion of causation but deny transitivity. For example, one might hold that counterfactual dependence is necessary as well as sufficient for causation. That does not yield especially plausible consequences — in particular, causation must be denied wherever redundancy is present (cf. Coady 2004). But it is a logical possibility. And there may be others.

Second, although a majority of contemporary philosophers probably endorse the view that we can and should separate out a concept of causation entirely devoid of principles of causal selection, few since Mill have argued for that conclusion, and a

significant minority disagree. Selection seems to be essential for causal judgments to serve many of the purposes for which we make them. Prediction, explanation, manipulation, and the attribution of moral and legal blame all involve selection. Moreover in at least some of these contexts it is hard to see selection as a secondary process. For example, if a driver swerves off the road and hits a pedestrian, it seems that we blame the driver for the injury *because* the driver caused it. The unselective view, however, has it that we say the driver caused the injury *because* we blame her for it.⁷ This is a complete reversal of the intuitive order. And, as Jonathan Schaffer puts it, “This is the sort of stable intuition that philosophers normally treat as data rather than rubbish” (Schaffer 2005: 343). The dim view of causal selection arises in part from the context-sensitivity of selective judgments; but that could just as well provide evidence for the view that causal judgments are context-sensitive (Menzies 2004, 2007), and that the relation that obtains when those judgments are true has places for more than just cause and effect (Maslen 2004, Schaffer 2005, Northcott 2008). The development of formal models for causal inference in statistics seems to employ a selective notion of causation: whether an event is a cause depends on whether it is counted as an endogenous or exogenous variable in a given model, and only then on whether it makes a difference of some kind within that model (see for example Pearl 2000, Halpern and Pearl 2005). Causal selection is contentious, complex and ill-understood: but it is not so obviously a caprice that, in denying that it has anything to do with causation itself, we can justify asserting causal transitivity also.

Third, adherence to an unselective notion of causation does not address all the counterexamples to transitivity. It even makes some of them harder to explain. So far we have been considering cases where transitivity appears to fail because the causal chain is somehow *too long*. Let us call cases of this sort *distance counterexamples* to transitivity. A robust stance against causal selection suggests (albeit vaguely) that we should also adopt a robust attitude towards distance counterexamples, denying that they are truly failures of transitivity, and instead explaining our reluctance to pronounce the events in question cause and effect by our interests, conventions, or other factors extraneous to the causal facts.

But spatiotemporal distance is not the only source of intuitive resistance to causal transitivity. Another is *double prevention*. Cases of double prevention are cases where c prevents some event d' , which would otherwise cause e , but c also causes d , which causes e . Thus d prevents c from preventing e : hence “double prevention”. In such cases, it seems that we have a strong intuitive disinclination to say that c causes e .

The following example illustrates this structure.

Dog-Bite Michael McDermott has occasion to detonate a bomb. The day before he does so, a dog bites off his right forefinger (c), preventing him from detonating the bomb with that finger (d'). So he detonates the bomb by pushing the button with his left forefinger (d). If the dog-bite (c) had not occurred, then the left-handed button-push (d) would not have occurred, and if that left-handed button-push (d) had not occurred, the bomb would not have detonated (e). If causation is transitive, and if the truth of each of these counterfactuals suffices for causation,

⁷Elsewhere I have argued that this is an untenable consequence in legal contexts (Broadbent 2009).

it follows that the dog-bite caused the explosion. Yet intuitively, the dog-bite does not cause the explosion. (Adapted from McDermott 1995: 531.)

This counterexample is not happily handled in the same way as distance counterexamples. It is not appealing to insist that the dog bite is a cause of the explosion, but not a salient one (cf. Lewis 2004: 98).⁸ All the obvious pragmatic requirements for salience are satisfied: it is spatiotemporally proximate, it is unusual, it is not the sort of information we would assume most audiences to possess already. Maybe the defender of transitivity can find other grounds on which to deny the intuition that the first event in the chain does not cause the last: but it is hard to see what they might be.

The only other option for the defender of transitivity is to deny that cases of double prevention exhibit true causal chains. One such strategy is to insist that the cases involve some illegitimate shift in level of description regarding the intermediary event. For example, for the purposes of picking out the cause of the explosion, the appropriate level of description is not “left-handed button push”, but simply “button push”. The left-handedness is irrelevant. And it is false that the dog bite caused the button push simpliciter. A more sophisticated implementation of this strategy is Laurie Paul’s suggestion that causation relates not events but aspects of events (Paul 2004). Thus the dog bite causes the button push to have the aspect of being left-handed, but this aspect does not cause any aspect of the explosion. Conversely the aspects of the button push that do cause aspects of the explosion (timing, sufficient force to connect the wires, etc.) are not caused by the dog bite, since they do not depend on which hand pushed the button.

Unfortunately this strategy does not handle all cases of double prevention. The trouble, as identified by Ned Hall, is that it requires d and d' to be reasonably described as the same event. But we can create double prevention cases where d and the event d' (which would have occurred had c not occurred) are not easily seen as the same event:

Suppose that after the dog-bite, the man does not push the button himself but orders an underling to do so. The relevant intuitions do not change: the dog-bite causes the order, and the order causes the explosion, but the dog-bite does not cause the explosion. The only way I can see to apply Paul’s observations is by way of a rather strained insistence that there is one event — call it a “making the button depressed” — which the dog-bite causes to have the aspect “being an order,” and which otherwise would have had the aspect “being a button-pushing.”

(Hall 2004a: 187)

Another, simpler case which is not amenable to the Paulian strategy is due to Hitchcock:

⁸Lewis argues that if we are happy to accept that causes can be redundant, we have no reason to aver in cases of double prevention. But this argument ignores our reason for accepting redundant causation. We have a strong intuition that at least some redundant causes are causes (e.g. Able’s throw in the previous example). If we have a strong intuition *not* to accept double prevention cases as cases of causation, then we *do* have a reason to aver in these cases: just as good a reason as we have for calling redundant causes causes. And most people do seem to be intuitively resistant to accepting that the first and last event in a double prevention case are cause and effect.

Boulder Christopher Hitchcock is hiking in the mountains when he spots a boulder bouncing towards him (*c*). He ducks (*d*) and survives (*e*). The boulder bouncing causes him to duck, and ducking causes him to survive. But the boulder bouncing does not cause him to survive.

Here, the alternative to ducking is presumably continuing to hike in an upright posture (*d'*). One might insist that the true cause of survival is the continued operation of biological systems such as respiration, and that ducking and walking upright are mere variants on this. But this does not help: aspects of the boulder bouncing (its direction, timing, velocity, size) cause the ducking-aspect of respiratory continuity; and the ducking-aspect of respiratory continuity causes many aspects of subsequent respiratory continuity. The heart will beat faster, the breathing will be accelerated, the memoirs will be enlivened, dinner parties improved, and so on.

Double prevention cases therefore provide very solid counterexamples to the claim that causation is transitive. Unlike distance counterexamples, it is not easy to see our resistance to calling *c* a cause of *e* in a double prevention case as a mere lack of salience. *c* is typically highly salient, because *c* is a threat to the occurrence of *e*. The defender of transitivity is forced to deny that double prevention consists in a chain of causes and effects. There is some intuitive plausibility to this denial: the chains in double prevention cases *are* a bit odd. But characterizing that oddness is difficult, as we have just seen.

What is more, an unselective approach to causation makes it even harder to deny that the links in the chain amount to causation. Take, for example, counterfactual dependence as an unselective sufficient condition for causation. In Boulder, there is counterfactual dependence of duck on boulder bouncing, and of survival on duck. There is likewise counterfactual dependence at the level of aspects: the ducking-aspect of respiratory continuity depends on aspects of the boulder bouncing, and many aspects of subsequent survival depend upon the ducking-aspect. If we adopted a more discriminating attitude towards causes, we might deny that the boulder bouncing caused the duck: we might say that the cause of the duck was Hitchcock's quick thinking, for example. But such discriminations are not available if we take counterfactual dependence as sufficient for causation.

Maybe the culprit here is not an unselective notion of causation, but more specifically the thesis that counterfactual dependence is sufficient for causation? Hall sees double prevention cases in this way: as conflicts between the transitivity thesis and the thesis that counterfactual dependence suffices for causation. His conclusion is that we should give up the latter claim. Though I have not argued it here, I agree with him about that: but my reasons are different. I think that counterfactual dependence is less selective than our causal concept (Broadbent 2008). Whereas Hall thinks that giving up the sufficiency of counterfactual dependence is the price we pay for preserving transitivity:

That causation is, necessarily, a transitive relation on events seems to many a bedrock datum, one of the few indisputable a priori insights we have into the workings of the concept. (Hall 2004a: 181)

But I do not share this insight. Moreover it seems that the *reason* the sufficiency thesis conflicts with the transitivity thesis is that it is highly unselective. The boulder rolling

is clearly part of the causal history, in the broadest sense, of the duck. There are nomic connections between the two events; energy is transferred (photons bouncing off the rock hit Hitchcock's eye); and so on. *If* we want to deny that the boulder rolling causes the duck, then we must find some principled way to discriminate among events that are not only counterfactually but also nomically and energetically connected to the effect in question, so as to honor some but not others with the title of "cause". If Hall's diagnosis is on the right lines, then the conflict is not between counterfactual dependence and transitivity: it is between the transitivity thesis and the thesis that causation is unselective.

In sum: causal selection is a curious and complex phenomenon. But adopting the view that causation is unselective does not provide a good justification for adopting the view that causation is transitive. The link is vague (first problem). The correct view of selection is not apparent (second problem). And at best, an unselective attitude towards causation handles only failure of transitivity due to spatiotemporal distance, where the distant cause may fail to be salient (third problem). It does not handle transitivity failure due to double prevention, where the cause is typically salient as a threat to the effect. And it makes it harder to deny that cases of double prevention amount to true causal chains, since an unselective attitude deprives the transitivity defender of one way to deny that each of the links is truly a case of causation.

4 Adding Up

Neither causal redundancy nor (scepticism about) causal selection provide strong reasons to think causation is transitive. The strongest reason, to my mind, comes from comparing causation to non-transitive relations. Consider, for example, touching. If c is touching d and d is touching e , then it is an open question whether c is touching e . The latter is logically independent, in the absence of any information about the locations of c , d and e . More specifically, if it is the case that c is touching e , that fact is intrinsic to c and e and does not have anything to do with the fact that each are touching d . Breaking one of the links, so that for example c does not touch d , does not necessarily mean that c no longer touches e .

Let us suppose, for the sake of argument, that the foregoing counterexamples to transitivity are accepted as showing that causation is non-transitive. Now consider a case where c causes d and d causes e , but where c also causes e . In other words, this is a causal chain that does not provide a counterexample to transitivity: for example, c is Able's throw, d is his rock hitting the bottle, and e is the bottle smashing. In this case it seems that the fact that c causes e is not independent of the intermediate causal links. Rather, it arises out of them. We would not say that Able's throw causes the bottle to smash if it did not cause his rock to hit the bottle, for example. (This is exactly why we deny that Baker, who also throws but misses, causes the smash.) Nor would we say Able's throw causes the smash if it hit the bottle but something else causes the smash — for example if the throw is feeble but the rock arrives at the same moment as, or just before, a meteorite.

Of course there could, in principle, be cases where c causes d , d causes e , and c causes e in a way that is independent of causing d . Perhaps c causes e by another

route. But not all cases are like that! There are many where c causes e by causing some d which causes e . In these cases, causal relations seem to “add up”, so that a series of causal relations give rise to one overarching one. If causation is transitive, that is exactly what you would expect. But if causation is non-transitive, then it is remarkable. Touching is not like that: the facts that c touches d and that d touches e never add up to, give rise to, or in any other way bear upon whether c touches e , except insofar as they constrain facts about the dimensions and locations of c and e , which are the determinants of whether c touches e .

What else, besides transitivity, could the “adding up” of causal relations consist in? What other kind of logical dependence of overarching relation on intermediaries is there? Any account of causal non-transitivity must answer this question. In my view, existing accounts do not do so adequately.

Suppose, for example, we impose a further condition on c causing e , such that c must raise the chance of e .⁹ Any such condition is likely to run into difficulties in cases of causal redundancy. But suppose further that some solution can be found.¹⁰ Then we have an account of *when* causal links “add up” to form an overarching one: when c also raises the chance of e . But we have no account of *how* or *why*. As with touching, it seems to be an independent matter whether c raises the chance of e . A defender of this approach could assert that it is *because* c raises the chance of d which raises the chance of e that c raises the chance of e . But what would that “because” mean? It does not denote logical necessity: it is already granted that there are other cases where causal links, and therefore (on this proposal) chance-raising, fail to add up to an overarching causal link, and therefore (on this proposal) chance-raising relation. Again, it is hard to see how this “because” could denote a causal link: surely the overarching causal relation is not *caused* by the links in the chain (even setting aside circularity worries), since it is natural to see the links as proper parts of the overall connection. No other interpretations of “because” suggest themselves. The assertion is not explanatory. If it is true at all, it states an aspect of the explanandum, and does not provide the explanans.

This problem afflicts several otherwise-promising solutions in the literature. For example, various authors have sought to extend the “contingencies” strategy for explaining causal redundancy (discussed in Section 2) to explain causal non-transitivity. Thus for example Hitchcock insists there must be an “active route” (Hitchcock 2001: 274) connecting cause to effect, where an active route is revealed by the counterfactual dependence of effect on cause, when the variables *not* belonging to that process are held fixed at their actual value. In Halpern and Pearl’s similar proposal, the contingencies revealing counterfactual dependence of effect on cause are not limited to holding variables fixed at their actual values (Halpern and Pearl 2005). Intuitively, Hitchcock’s solution says that an active route connects c and e when e counterfactually depends upon c holding intermediaries not belonging to the causal process fixed at their actual value, while Halpern and Pearl require only that there be *some* assignment of values of the intermediaries under which e counterfactually depends on c . On both accounts, if an active causal route leads from c to e then c causes e . And on both accounts,

⁹I am grateful to an anonymous referee for this suggestion.

¹⁰So that, where c_1 causes e and c_2 is a preempted back-up, there continues to be some sense in which c_1 raises the chance of e even where c_2 make e just as probable in c_1 ’s absence. For one proposal see Dowe 2004.

cases of double prevention fail this test. In Boulder, there are no intermediaries that are not on the causal pathway, so nothing to hold fixed (in Hitchcock's account) or to vary (in Halpern and Pearl's). Therefore the boulder rolling does not cause Hitchcock's survival.

This is a promising solution. But once again, it does nothing to explain why causation "adds up". Indeed, on this approach, whether a causal chain adds up to an overarching causal relation seems more extrinsic than ever. Even if these accounts achieve extensional accuracy,¹¹ they do not make clear how the links in the chain are related to the overarching causal relation, when there is one. Yet it is intuitively clear that there is a close, even a constitutive connection here.

If causation is not transitive then we need to do more than just provide an extensionally accurate account of our causal judgments. An account of causal non-transitivity needs to do more than just get the cases right. It also needs to explain how causation can be non-transitive and yet, on occasion, display additive features characteristic of a transitive relation. Extant accounts have not adequately addressed this explanatory task.

5 Two Kinds of Difference-Making

One obvious, if banal, thing to say about double prevention cases is that c does not make any difference to e . Making a difference is a vague notion. The solution I propose requires us to distinguish two ways in which this vagueness might be resolved. The first is this:

Cause-led difference-making A cause makes a difference in that had it been different or absent then its effect would have been different or absent.

The clearest expression of this idea remains David Lewis's counterfactual analysis of causation, which starts with this idea:

We think of a cause as something that makes a difference, and the difference it makes must be a difference from what would have happened without it. (Lewis 1973a: 160)

Many other approaches are variations on this theme. Thus, as we have already seen, Halpern and Pearl specify that what counts is not what would have happened had the effect simply failed to occur, but had it failed to occur under certain contingencies (Halpern and Pearl 2005); and Hitchcock has developed accounts of various causal phenomena exploiting the causal modelling framework, and the notion of difference-making that it employs (see e.g. Hitchcock 2001, 2007). Menzies has proposed adding the notion of a "default value", relative to which causes must vary in order to count as difference-makers (Menzies 2004, 2007). Schaffer doubts that a non-circular conceptual analysis of causation can be had (Schaffer 2004), but proposes that the causal relation is four-place, with a contrast for both cause and effect; and that a decent test

¹¹One obvious criticism is that on both accounts the truth of a causal claim is relative to a model, and principles of model selection are very difficult to articulate, as these authors all acknowledge.

for causation is a variant of counterfactual dependence of effect on cause, such that the counterfactual non-occurrence of cause and effect consists in the occurrence of events specified explicitly or implicitly by our four-place contrastive causal claims (Schaffer 2005: 328-9). Carolina Sartorio analyses causation explicitly in terms of another variation on the same theme of difference-making: the variation being that to qualify as the right kind of difference-making for causation, the absence of the cause must not itself cause the effect (Sartorio 2005: 75). And the central component of James Woodward's "manipulationist" account is that causal explanations answer "what-if-things-had-been-different" questions, where it is clear that the things in question are causes (Woodward 2003: 11).

Cause-led difference-making suffers from obvious defects as an analysis of the sort of difference causes make. Sometimes, in cases of redundancy, causes do not seem to make this sort of difference. Moreover, there are events which make this sort of difference but which we discriminate against, selectively treating just a special few as causes for most practical purposes. Nor is cause-led difference-making much use with our present task, that is, explaining causal non-transitivity. If we require that causes make this sort of difference to their effects, then we reach the intuitively correct conclusion that double prevention is not causation — but we run into difficulties with redundant causes. Moreover distance counterexamples to transitivity may satisfy this condition, since events in the distant past which we do not treat as causes for most purposes may nevertheless make this sort of difference (e.g. my writing these words counterfactually depends on the Big Bang). And because this sort of difference-making is not transitive, we have no explanation of how causal chains can "add up" to yield an overarching causal connection. This problem persists even if more sophisticated variants of this approach (e.g. Hitchcock's, or Halpern and Pearl's) can achieve extensional accuracy in both double prevention and redundancy cases.

But here is another, non-equivalent kind of difference-making:

Effect-led difference-making A cause makes a difference in that it is a difference between the effect being as it is and the effect being different or absent.

Effect-led difference-making is familiar from the literature on contrastive explanation. Whether a given event is a good causal explanation of a contrast between a fact and a foil seems to depend upon whether it makes a difference in the effect-led sense. Making a difference in the cause-led sense is not enough. For example, if you ask why I arrive late rather than on time, mentioning the presence of oxygen will not be a good answer, even though oxygen makes a cause-led difference (without oxygen I would not have arrived late, or at all). The reason it is unhelpful to mention the oxygen is that it is present in both the fact (the actual late arrival) and the foil with which it is contrasted (the on-time arrival). It does not make an effect-led difference. Whereas my oversleeping is a difference between fact and foil.¹²

¹²This sort of difference-making may not be sufficient for explanatory relevance — there remain many irrelevant differences between fact and foil — but their number is considerably reduced. Moreover, as Lipton points out, this approach helps explain why explanation appears to have both objective and subjective aspects. On this analysis, it is up to us what contrast to explain, but then it is an objective matter whether a given event is a difference between the fact and foil we have selected (Lipton 1990).

Several theorists have proposed necessary conditions on explaining contrasts that employ the effect-led notion of difference-making. David Lewis's account of contrastive explanation endorses a counterfactual version of this idea, whereby an explanatory cause is one that would not have featured in the counterfactual causal history of the foil, had the foil occurred Lewis [1986a]. Peter Lipton's account tends towards the actualist: an explanatory cause is one an event in the causal history of the fact, lacking a corresponding event in the actual history of the failure of the foil to occur (Lipton 1990, 2004, Ch 3 — but see Lipton 1993 for a fuller discussion of the actual/counterfactual issue). Yet effect-led difference-making has not attracted much attention at all in the analysis of causation itself.

My suggestion is that effect-led difference-making is a necessary condition, not just on contrastive causal explanation, but on causation itself. To be a cause, an event must make a difference to its effect in the effect-led sense.

There may be more than one way to put flesh on the bones of effect-led difference-making, but the simplest way to cash the idea out is with a counterfactual:

The Reverse Counterfactual If c causes e then $\neg Oe \Box \rightarrow \neg Oc$.

This is the familiar Lewisian counterfactual, reversed, and proposed as a necessary (not a sufficient) condition on causation.¹³

Obviously, this is a backtracking counterfactual, and therefore on Lewis's semantics false under standard resolutions of vagueness. I cannot provide an alternative to Lewis's semantics here. Nonetheless there are reasons to countenance the proposal.

First, this is just a way of cashing out the notion of effect-led difference-making. If it helps with causal transitivity then we will have discovered something interesting, even if the particular vehicle we used to get there is a bit crude.

Second, there are good reasons to suppose that the strong asymmetry of counterfactual dependence proposed by Lewis is incorrect, some of which are discussed in (Broadbent 2007). It is doubtful that Lewis's account of the asymmetry of counterfactual dependence yields that asymmetry (Bennett 2001, Elga 2000). And it is doubtful whether the asymmetry exists in the strong form he asserts in the first place. Even Lewis's own account of contrastive explanation appears to require backtracking, since it requires us to ask what would have been the case in the past of the counterfactual foil to the actual effect being explained (Lipton 1990). Any counterfactual sensitivity condition on evidence will require that backtracking counterfactuals hold whenever causes provide evidence for their effects (Lipton 2000; Broadbent 2007: 172). Dan Hausman has argued that we sometimes backtrack even to assess foretrackers: for example, engineers modelling the consequences of a burst pipe in a nuclear plant would not simply suppose that the pipe had burst due to a small miracle, but would model the consequences given various plausible scenarios in which the pipe would have burst — terrorist attack, earthquake, and so on (Hausman 1998: 121-2).¹⁴

¹³If the Reverse Counterfactual is a necessary condition on causation then counterfactual dependence cannot be a sufficient condition on causation, without yielding the counterintuitive result that effects cause their causes whenever redundancy is absent and Lewis's counterfactual is true.

¹⁴Hausman himself does not accept that backtracking counterfactuals have determinate truth-values, however.

Third, if what is needed is some assurance that we actually do assign truth values to backtracking counterfactuals, then an epistemology of counterfactuals will be as useful or more useful than a semantic theory. Timothy Williamson proposes that counterfactual reasoning is akin to simulation, employing offline all the same cognitive faculties we employ online: the sort of abilities we use to predict where a rock bouncing down a hill will land. We work out where the rock would land by employing these same faculties – whatever they are — offline, under some counterfactual supposition about where the rock began. This “imaginative exercise is radically informed and disciplined by your perception of the rock and the slope and your sense of how nature works” (Williamson 2007: 143). But we can equally work out where the rock must have come from. That is just as useful an ability, and just as obviously one we possess, as being able to predict where it will land. And this imaginative exercise is just as radically constrained by our sense of how the world works as the forward-looking one. This suggests that we can work out where the rock *would* have come from, had it landed somewhere else. If anything like this account of counterfactual reasoning is correct, then there is every reason to suppose that we have backtracking abilities as well as foretracking ones: and therefore that we can assign truth-values to backtracking counterfactuals as well as foretracking ones.

The only caveat I add here is that I do not claim that the Reverse Counterfactual is *intuitively* obvious; often it will sound very strange. I suggest that it is true nonetheless, although — as Lewis remarks — we are inclined to use other grammatical constructions to express backtrackers (cf. Lewis 1979: 34–35). In my view, that is not because they are false, but because the subjunctive mood has other conversational implications, perhaps most obviously causal ones.

6 Solution

How does effect-led difference-making, in the form of the Reverse Counterfactual, help understand causal non-transitivity?

First, it appears to get the extension correct in double prevention cases. Supposing e not to have occurred, in each case, means supposing that something brought about its non-occurrence. The examples already feature a “preventer” — an event c which would, left to its own devices, prevent the effect e . So there is a salient and obvious way to suppose that the effect would not have occurred: simply suppose away the second preventer, d , which prevents c from preventing e . This means it is generally false that $\neg Oe \Box \rightarrow \neg Oc$.

For example, in Boulder, had Hitchcock not survived ($\neg Oe$), that would presumably be because he failed to duck ($\neg Od$). Thus the duck satisfies the necessary condition imposed by the Reverse Counterfactual for causing the survival. But surely it is false that, had Hitchcock not survived ($\neg Oe$), the boulder would not have rolled ($\neg Oc$). What, then, would have killed him?¹⁵ Similarly, in Dog Bite, it seems reasonable to

¹⁵Indeed at one point Hitchcock almost stumbles into asserting the Reverse Counterfactual: “Intuitively, the falling boulder does not save Hiker’s life because without it, Hiker’s life would not have been endangered in the first place. This is just what is indicated by the absence of an active route from [the boulder falling] to [Hiker surviving]: there is no scenario in which the boulder does not fall and Hiker does not survive.”

suppose that, had the explosion not occurred ($\neg Oe$) then the button would not have been pushed ($\neg Od$). But there is no sense in the suggestion that, had the explosion not occurred ($\neg Oe$), then the dog bite would not have occurred ($\neg Oc$). There is simply no reason to say so.

Second, the Reverse Counterfactual does more than get the extension right: it also explains why, in cases of double prevention, the causal links do not “add up”. Counterfactuals are not transitive, but a valid substitute for transitivity is this:

$$B \Box \rightarrow A$$

$$A \Box \rightarrow B$$

$$B \Box \rightarrow C$$

$$A \Box \rightarrow C$$

(Lewis 1973b: 17)

Now consider these counterfactuals:

$$(1) \quad \neg Od \Box \rightarrow \neg Oe$$

$$(2) \quad \neg Oe \Box \rightarrow \neg Od$$

$$(3) \quad \neg Od \Box \rightarrow \neg Oc$$

$$(4) \quad \neg Oe \Box \rightarrow \neg Oc$$

If these counterfactuals were all true, then it would follow that $\neg Oe \Box \rightarrow \neg Oc$. Yet in cases of double prevention these counterfactuals are not true together.

Consider Dog Bite. It does not seem particularly reasonable to suppose that if the push had not occurred ($\neg Od$), the bite would not have occurred ($\neg Oc$). — Why on earth not? So (3) seems to be false. To make (3) plausible, we need to specify that the button push d be *left-handed*. Had *that* not occurred, let us grant, then the dog would not have bitten. But now (1) becomes implausible. Had the left-fingered button push not occurred ($\neg Od$), and (as specified in (3)) the dog not bitten ($\neg Oc$), then the explosion might still have occurred ($(\neg Oc \& \neg Od) \diamond \rightarrow Oe$): because the right forefinger might have been used.

Similar remarks apply to Boulder. It is hard to see how to get the appropriate combination of counterfactuals to be true together. Had Hitchcock not ducked ($\neg Od$), would the boulder have rolled (Oc)? If we say yes, then (3) is false. If we say no, then (1) is false — if Hitchcock fails to duck and the boulder does not roll, then surely he would survive.

Consider another example:

(Hitchcock 2001: 295) This is not far from saying that, if Hiker had not survived, the boulder would still have fallen: but Hiker would have omitted to duck.

Hospital generator There is a power-cut c , the generator comes on d , and the hospital lights are on e . The power-cut causes the generator to come on and the generator coming on causes the lights to be on. Yet the power-cut does not cause the lights to be on.

This is another case of double prevention, and again, it is hard to see how (1)–(3) can all be satisfied together. Had the generator not started, would the grid have failed? If yes, then (3) is false. If no, then (1) is false. And so on for the other examples. This is not a purely technical point. It reflects the fact that, in cases of double prevention, we are not prepared to simultaneously regard the intermediate event, d , as caused by c and as counterfactually necessary for e .

Third, and conversely, the Reverse Counterfactual helps us to understand causal non-transitivity by providing a way for intermediate causal links to “add up” to one big one. There seems to be some logical relation between the links in a causal chain and the overarching causal relation, when there is an overarching relation. I argued in Section 4 that otherwise promising extant accounts are inadequate because they do not explain this fact. However, the Reverse Counterfactual does provide a logical connection between the links in the chain and the overarching causal relation. When conditions (1)–(3) are satisfied, the Reverse Counterfactual follows. When they are violated, it does not.¹⁶

It might be objected that I have proposed the Reverse Counterfactual as a necessary but not a sufficient condition on causation, and so the fact that the Reverse Counterfactual is implied in certain cases cannot be taken to explain why causation seems to be implied in those cases. I acknowledge that the explanation is only partial. But it is of some value nevertheless. It tells us why breaking one of the links in a causal chain that adds up to an overarching causal relation also entails breaking the overarching causal relation: because then there is nothing to ensure that the necessary condition imposed by the Reverse Counterfactual is met. Moreover the fact that the truth of the Reverse Counterfactual seems to vary with our intuitions on transitivity also suggests that whatever ingredients we would need to add to achieve a sufficient condition, they are not relevant to the question of transitivity.

As well as handling double prevention, this explanation also suggests a diagnosis of the failure of causal transitivity with distance. If we take longer chains, the likelihood that the Reverse Counterfactual will be implied across the length of the chain drops, since it becomes less likely that the requisite conditions will be satisfied. For them to be satisfied, the chain must be such that, first, the Reverse Counterfactual is satisfied at every step; and second, there is foretracking counterfactual dependence of effect on cause at (almost) every step. This second requirement means that there must be no (or little) redundancy. So perhaps the reason that we are reluctant to see my birth as a cause of my writing these words, for example, is that we worry that the long

¹⁶Unless another valid substitute can be found. On inspection it turns out that the only other hopeful substitutes yield the same consequences that this one does. For example, consider $A \square \rightarrow B$, $(A \& B) \square \rightarrow C$, hence $A \square \rightarrow C$. In Boulder, it seems reasonable to say: if Hitchcock had not survived, he would not have ducked. But it does not seem reasonable to say: if Hiker had neither survived nor ducked, the boulder would not have fallen. For the obvious scenario where he dies, and also fails to duck, is one where the boulder hits him. To suppose that the boulder did not fall but that something else entered the picture to kill him would be entirely gratuitous.

and largely unknown chain may fail one of these conditions. In particular, we might wonder if it includes some redundancy. In fact, I was educated in Dullsville; but had I been educated in Boretown, it might not have made a difference. Or perhaps there is not a chain of causes and effects — in the strong, selective sense which Reverse Counterfactual suggests (Broadbent 2008). Perhaps my eating breakfast is not naturally seen as the cause of my writing these words, for example.

It would be nice to check the correctness of this proposal with a little thought experiment. Suppose someone provided a very long chain of events, and it was convincing that each was a cause of the next (implying, on my account, the satisfaction of the Reverse Counterfactual) and moreover that there was no causal redundancy. Then, if I am right, we would be inclined to accept that the first event also caused the last. And indeed there is a proverb which does exactly this:

Horse shoe For want of a nail the shoe was lost; for want of a shoe the horse was lost;
for want of a horse the rider was lost; for want of a rider the battle was lost; for
want of a battle the kingdom was lost — and all for the want of a horseshoe nail.

The psychological effect of this rhyme is precisely to convince us, or at least suggest, that the lack of a nail caused the loss of the kingdom.

In general, however, it is extremely hard to find examples of this kind. The more events we include, and the longer the chain gets, the easier it is to wonder whether one of these events might have come about anyway, by some other means. Or we might wonder suspect that one of the links is not a really causal, in the strong selective sense that we revert to when we are not doing philosophy, in which case the Reverse Counterfactual will not hold at that step (Broadbent 2008). In either case, there is no guarantee that the Reverse Counterfactual will hold between the first event in the chain and the last. I have suggested that this is why we are not compelled to take the first event in the chain as the cause of the last.

On this analysis, then, it appears that failure of transitivity occurs for the same reason in cases of double prevention as it does in cases of distance. The difference is epistemic. In cases of double prevention, we *know* that the intervening links necessary to render the Reverse Counterfactual true are lacking: we can see that there is redundancy, or that a putative causal link in the chain has a fishy, philosophical smell, which we would normally discriminate against. Whereas in cases where transitivity fails due to distance, we merely suspect that there might be some breakdown in the long and largely unknown chain.

In either case, the explanation is the same. Causes must make a difference to their effects, where that is an effect-led difference. This sort of difference-making is implied when certain patterns of difference-making are present along the causal chain, explaining how causal relations can be additive without being transitive. I have crudely characterised this sort of difference-making with the Reverse Counterfactual; but there may be other ways to do it.¹⁷

¹⁷I am grateful to Arif Ahmed, Alexander Bird, Peter Lipton and the anonymous referees at this journal for helpful comments, and to the PHG Foundation and the Brocher Foundation for supporting this research.

References

- Jonathan Bennett. On forward and backward counterfactual conditionals. In G. Preyer and F. Siebelt, editors, *Reality and Humean Supervenience: Essays on the Philosophy of David Lewis*, pages 177–202. Rowman and Littlefield, Maryland, 2001.
- Alex Broadbent. The difference between cause and condition. *Proceedings of the Aristotelian Society*, 108:355–364, 2008.
- Alex Broadbent. Fact and law in the causal inquiry. *Legal Theory*, 15:173–191, 2009.
- Alex Broadbent. Reversing the counterfactual analysis of causation. *International Journal of Philosophical Studies*, 15:169–189, 2007.
- David Coady. Preempting preemption. In J. Collins and L. A Paul, editors, *Causation and Counterfactuals*, pages 325–340. MIT Press, 2004.
- Phil Dowe. Chance-lowering causes. In P. Dowe and P. Noordhof, editors, *Cause and Chance: Causation in an Indeterministic World*, pages 28–38. Routledge, London, 2004.
- Adam Elga. Statistical mechanics and the asymmetry of counterfactual dependence. *Philosophy of Science (Proceedings)*, 68:S313–S324, 2000.
- Ned Hall. Causation and the price of transitivity. In J. Collins, N. Hall, and L. A Paul, editors, *Causation and Counterfactuals*, pages 181–204. MIT Press, Cambridge, Massachusetts, 2004a.
- Ned Hall. Two concepts of causation. In J. Collins, N. Hall, and L. A Paul, editors, *Causation and Counterfactuals*, pages 225–276. MIT Press, Cambridge, Massachusetts, 2004b.
- Joseph Y. Halpern and Judea Pearl. Causes and explanations: A Structural-Model approach. part i: Causes. *British Journal of the Philosophy of Science*, 56:843–887, 2005.
- Daniel Hausman. *Causal Asymmetries*. Cambridge University Press, Cambridge, 1998.
- Christopher Hitchcock. The intransitivity of causation revealed in equations and graphs. *The Journal of Philosophy*, 98:273–299, 2001.
- Christopher Hitchcock. Prevention, preemption, and the principle of sufficient reason. *Philosophical Review*, 116(4):495–532, 2007.
- David Lewis. Causal explanation. In *Philosophical Papers, Volume II*, pages 214–241. Oxford University Press, Oxford, 1986a.
- David Lewis. Causation. *Journal of Philosophy*, 70:556–567, 1973a. Page numbers refer to (Lewis 1986).

- David Lewis. Causation as influence. In J. Collins, N. Hall, and L. A Paul, editors, *Causation and Counterfactuals*, pages 75–106. MIT Press, Cambridge, Massachusetts, 2004.
- David Lewis. Counterfactual dependence and time’s arrow. *Noûs*, 13:455–476, 1979. Page numbers refer to (Lewis 1986).
- David Lewis. Counterfactuals and comparative possibility. *Journal of Philosophical Logic*, 2:418–446, 1973b. Page numbers refer to (Lewis 1986).
- David Lewis. *Counterfactuals*. Harvard University Press, Cambridge, Massachusetts, 1973c.
- David Lewis. *Philosophical Papers, Volume II*. Oxford University Press, Oxford, 1986b.
- Peter Lipton. Making a difference. *Philosophica*, 51(1):39–54, 1993.
- Peter Lipton. Contrastive explanation. In D. Knowles, editor, *Explanation and its Limits*, pages 246–266. Cambridge University Press, Cambridge, 1990.
- Peter Lipton. *Inference to the Best Explanation*. Routledge, London and New York, 2 edition, 2004.
- Peter Lipton. Tracking track records. *Proceedings of the Aristotelian Society — Supplementary Volume*, 74(1):179–205, 2000.
- Cei Maslen. Causes, contrasts, and the non-transitivity of causation. In John Collins, Ned Hall, and L. A Paul, editors, *Causation and Counterfactuals*. MIT Press, Cambridge, Massachusetts, 2004.
- Michael McDermott. Redundant causation. *The British Journal for the Philosophy of Science*, 46(4):523–544, 1995.
- Hugh Mellor. *The Facts of Causation*. Routledge, London and New York, 1995.
- Hugh Mellor. For facts as causes and effects. In J. Collins, N. Hall, and L. A Paul, editors, *Causation and Counterfactuals*. MIT Press, Cambridge, Massachusetts, 2004.
- Peter Menzies. Causation in context. pages 191–223. Oxford University Press, Oxford, 2007.
- Peter Menzies. Difference-making in context. In J. Collins, N. Hall, and L. A Paul, editors, *Causation and Counterfactuals*, pages 139–180. MIT Press, Cambridge, Massachusetts, 2004.
- Robert Northcott. Causation and contrast classes. *Philosophical Studies*, 139:111–123, 2008.
- Robert Nozick. *Philosophical Explanations*. Harvard University Press, Cambridge, Massachusetts, 1981.

- L. A Paul. Aspect causation. In J. Collins, N. Hall, and L. A Paul, editors, *Causation and Counterfactuals*, pages 205–224. MIT Press, Cambridge, Massachusetts, 2004.
- Judea Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, Cambridge, 2000.
- Carolina Sartorio. Causes as Difference-Makers. *Philosophical Studies*, 123:71–96, 2005.
- Jonathan Schaffer. Contrastive causation. *Philosophical Review*, 114(3):297–328, 2005.
- Jonathan Schaffer. Counterfactuals, causal independence and conceptual circularity. *Analysis*, 64(4):299–309, 2004.
- Timothy Williamson. *The philosophy of philosophy*. Blackwell, Oxford, 2007.
- James Woodward. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, Oxford, 2003.