

The New Riddle of Causation

Alex Broadbent

December 11, 2007

Abstract

We commonly distinguish causes from mere conditions, for example by saying that the strike caused the match to light but by failing to mention the presence of oxygen. Philosophers from Mill to Lewis have dismissed this common practice as irrelevant to the philosophical analysis of causation. In this paper, however, I argue that causal selection poses a puzzle of just the same form as Hume's sceptical challenge to the notion of necessary connection. I then propose a solution in terms of a simple counterfactual.

1 Causes and Conditions

Suppose I strike a match, and it lights. In the present circumstances, sitting at my desk in a warm dry room, I would usually say that striking the match caused it to light. I would not usually say that the presence of oxygen caused it to light. This is despite my awareness that both strike and oxygen are required, in the circumstances, for the flame. I admit, in other words, that both these counterfactuals are true:

- If I had not struck the match, it would not have lit;
- If there had been no oxygen present, the match would not have lit.

You may not think much follows from the fact that I do not normally *say* the oxygen caused the match to light. Just because I don't say so, it doesn't follow that the oxygen is not in fact a cause — it could just be that, sitting in my warm, dry room, I have rather come to take the presence of oxygen for granted. You might further point out that, in a different context, the oxygen is mentioned as the cause. Hart and Honore give the example of a fire which starts when air leaks into a delicate manufacturing process normally conducted in a vacuum, contending that in that case we do say that the presence of oxygen caused the fire [Hart and Honore, 1985, p10]. You might suggest that this is because the oxygen is a cause, even in the warm, dry room where I take the presence of oxygen for granted. My lack of appreciation for the presence of oxygen does not, you might argue, bear on its causal powers.

That indeed appears to be the majority view. Mill says:

The real Cause, is the whole of these antecedents; and we have, philosophically speaking, no right to give the name of cause to one of them, exclusively of the others. [Mill, 1887, p237]

David Lewis has a different view, not taking the cause to be *all* of the antecedents, but taking *each* of them individually to be a cause:

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, p162]

Although Lewis’s view is different from Mill’s, the two writers agree that, vulgar talk notwithstanding, a philosophical account of causation ought not (or at least need not) explain our commonplace discriminatory practices.¹

Let us say that a *condition* for an actual event e is an actual and distinct event c such that $\sim O(c) > \sim O(e)$ (if c had not occurred, then e would not have occurred). In this special sense of “condition”, to say that c is a condition of e is to say that, in the circumstances, c is counterfactually necessary for e . (The term “condition” is of course already in use in established phrases like “necessary condition” and so forth; our conditions are belong to a subspecies — specifically, they are counterfactually necessary conditions. It is my job to make sure no confusion arises concerning whether my intended meaning is the ordinary, general sense or the special sense just defined.)

The cause/condition distinction is powerful, because it can be used to capture challenges to both the necessary and the sufficient components of a counterfactual analysis of causation. Not all causes are conditions, because of cases of *causal redundancy* (a generic term I use to cover preemption, symmetric overdetermination, and any other case where the cause fails to be counterfactually necessary for the effect). For this reason, Lewis did not suggest that the simple counterfactual which we used to characterise conditions is always true of causes: it is not necessary for causation that the effect depend upon the cause. Almost everybody else agreed with him (one exception is Coady [2004]), and the hunt was on to discover a necessary condition for causation.

But Lewis successfully convinced the philosophical world that a sufficient condition for causation had been provided in this simple counterfactual.² He

¹Jonathan Schaffer also juxtaposes Mill and Lewis, although he picks a different passage from Mill [Schaffer, 2005, pp312–313].

²In fact he thought that counterfactual *dependence* sufficed for causation; dependence further requires that $O(c) > O(e)$. However, Lewis requires that causes and effects must be actual and distinct events. On Lewis’s semantics it follows from the fact that c and e occur that $O(c) > O(e)$. So once the other actuality and distinctness requirements are met, whether some putative effect e counterfactually depends on some putative cause c turn solely

considered two sorts of challenge to sufficiency, asking whether causes also depended on their effects, and he asked whether effects of a common cause depended on each other. In both cases the answer was blunt:

The proper solution to both problems, I think, is flatly to deny the counterfactuals that cause the trouble. [Lewis, 1973a, p170]

This is perhaps rather surprising, especially as regards effects of a common cause. There are many measurement instruments and biological systems which work by tracking the cause of the phenomenon they are intended to measure, rather than that phenomenon itself. It is a consequence of Lewis's view that counterfactual reliability is unachievable for any such system. You consult your speedometer as you approach the speed camera, and are satisfied that you are on the right side of the law. You might be tempted to assume that if you were going faster, your speedometer would tell you so. But on Lewis's view that is false, for most actual speedometers; for they do not measure your speed, but the rate at which your wheels are rotating. It is possible to jack up your car so that the speedometer shows a high speed when you are in fact stationary. In other words, your speed does not cause your speedometer to read what it does. And so as you approach the speed camera and relate your musings to Lewis, he responds that if you had been going faster, the speedometer would have read just what it actually does: it would have been broken. Not only would you have got a speeding ticket, you would also have had to take your car to the garage: so it was just as well you slowed down.

I have sometimes encountered some resistance to the claim that Lewis holds counterfactual dependence to suffice for causation. But he does. It is clear in his theory of events, where he considers that the prospect of counterfactual dependence between events which are not related as cause to effect is the decisive reason to rule out both extrinsic and disjunctive events [Lewis, 1986b, pp262–269]. Moreover, through all his struggles to arrive at a satisfactory necessary component for his analysis of causation, he made it clear that counterfactual dependence still provided a sufficient condition: “We have a sufficient, but not a necessary, condition for causation” [Lewis, 2004, p78]. Where c and e are actual and distinct events, if $\sim O(c) > \sim O(e)$ then c causes e . — Provided, of course, that the counterfactual is “of the right sort”, not one which runs from effect to cause, or from effect to sibling effect.

Lewis's motive for denying that causes depend upon effects or that effects depend upon each other is clear, then: but what is his argument? He provides an argument in discussion of a system which is extremely similar to my speedometer: he uses a barometer. The only reason, he says, that we think

on whether $\sim O(c) > \sim O(e)$. Hence the lack of attention to $O(c) > O(e)$, here and throughout the literature. Lewis discusses the point [Lewis, 1973a, pp165–167], and in places assumes that counterfactual dependence is equivalent just to the counterfactual “ $\sim O(c) > \sim O(e)$ ” [Lewis, 2004, p78]. I shall accept without dispute Lewis's conception of the causal relata as actual and distinct events, and the aspect of his semantics which makes counterfactuals with true antecedents and consequents automatically true (that aspect is the Centering Assumption [Lewis, 1973b, pp26–31]).

the barometer reading depends on the weather, is that we reason back from the reading to its cause, atmospheric pressure, and then forward again to the weather. We reason that if the reading had been different, the pressure would have been different, and so the weather would have been different too. But the first step backtracks: the atmospheric pressure falls before the barometer falls. Lewis argues extensively that (in the ordinary context which determines the truth-value of his causal counterfactuals) backtracking counterfactuals are usually false.³ So we should not be tempted to think that causes depend upon their effects, since that involves backtracking reasoning which is, in normal circumstances, fallacious. And nor should we suppose that effects of a common cause depend upon each other, since that would also involve supposing that causes depend upon their effects, and thus backtracking reasoning.

A similar point applies to simple cause-effect pairs. Perhaps it is more usual to treat the effect as a reliable sign of the cause, but causes can also be reliable signs of their effects. Lewis's position prevents any counterfactual characterisation of this reliability, just as it prevents the counterfactual characterisation of the reliability of effects as signs of other effects.⁴ This is strange, given the close connection described by Hume between causation and inductive inference. To analyse causation but not induction with counterfactuals is to adopt a strange stance [a point also made in Broadbent, 2007, p172].

This is not the place to go into Lewis's reasons for denying backtracking counterfactuals. Note, however, how the choice of example affects the plausibility of the argument. The fluidity and imprecision of the mechanism, the time-lag, even the antiquity of the instrument in question, all make the argument much more plausible in the case of the barometer than it is in the case of the speedometer. Notice too that in the case of the speedometer, we could if we wished measure the speed directly (whereas we cannot predict the weather directly). If we did, then the reading would depend upon the speed. This allows us to bring out another curious feature of Lewis's view: whether or not the speedometer reading is counterfactually reliable depends upon whether it measures speed directly, or measures a cause of the speed. But the reliability of the speedometer is not affected by this distinction between mechanisms, except in the unusual circumstance that the car is up on jacks.

My purpose in this essay is to discuss a challenge to the sufficiency for causation of Lewis's counterfactual, which is different from the two just mentioned: namely, the *problem of causal selection*. This can also be expressed in terms

³I have sometimes encountered the view that what Lewis meant by a backtracker was a counterfactual running counter to the causal direction, i.e. from effect to cause. This is a misconception. Lewis clearly meant backtracking as a temporal notion. The prevalence of temporally foretracking counterfactuals, which Lewis argues is a contingent fact about our world, combines with the Lewis's counterfactual analysis of causation to explain the prevalence of forward causation in our world [Lewis, 1979]. If backtracking were meant as a causal rather than a temporal notion then no such explanation of the temporal direction of causation would be available.

⁴David Hume explicitly countenances inferences both ways between cause and effect, as well as inferences from one effect to another: "Heat and light are collateral effects of fire, and the one effect may be justly inferred from the other" [Hume, 1748, p73].

of the cause/condition distinction. The problem is that we do not normally treat all conditions as causes. In the match-strike example, the presence of oxygen is treated as a condition, but not a cause, of the flame. According to the Mill-Lewis tradition, this is because we are whimsical creatures with restricted interests; the distinction is obviously a feature of our treatment of causes, not of causes themselves, and therefore it is of no interest to you if you are interested in causation in the objects. But I shall argue that the problem is much deeper than that.

More precisely, I want to argue:

- (1) Causal selection is ubiquitous in our causal judgements; but
- (2) Causal selection is apparently without basis in the objects; and
- (3) (1) and (2) are in tension and generate a deep philosophical problem.

I shall argue for (1) in §2, and for (2) and (3) in §3, where I will explain why the problem of selection is a “new riddle”. In §4 I shall sketch my best effort at a solution for the new riddle.

2 Causal Selection is Ubiquitous

Jonathan Schaffer gives three reasons to reject the Mill-Lewis line on causal selection. First, he says, selection is predictable. When I strike the match and it lights, almost everybody selects the match-strike as the cause, at least in the context of a warm, dry room. As Schaffer says, “This is the sort of stable intuition that philosophers normally treat as data rather than rubbish” [Schaffer, 2005, p313]. Second, Schaffer points out the causal selection is central to our moral and legal practices, and moreover that it is treated in the courts as an objective question and not as a mere caprice. And third, Schaffer argues that selection is an “inseparable part of our causal concept” [Schaffer, 2005, p314]. Elsewhere [Broadbent, 2008] I develop arguments along the lines of Schaffer’s first two reasons, arguing that extant accounts of causal selection which take the Mill-Lewis as a premise are inadequate to explain the predictability of selection in any context, and the role of selection in moral and legal contexts. In this paper, I shall focus on the third reason Schaffer gives for doubting the Mill-Lewis line. Although I personally find Schaffer’s first two reasons compelling, others may not. It remains open to hold that the explanation of causal selection offered by Schaffer, me, or anyone else, is not sufficiently attractive to motivate dropping the Mill-Lewis view of causal selection. Philosophers are used to biting the bullets of intuition, so ignoring the predictability of selection will not be too difficult. And the fact that we treat selection as objective in moral and legal contexts is an inconclusive argument because it is so unclear whether the subject of moral and legal discourse is itself objective. Maybe causal selective judgements are just another way of trying to project our norms onto an amoral world. An argument developing these themes may be suggestive, but it will not be demonstrative.

The alternative argument which I develop in this paper is bolder, and the first respect in which it is bolder is this: I suggest that causal selection is not merely a common feature of causal judgements; I suggest it is an ubiquitous feature. In this section, I argue that every causal judgement which has ever been or will ever be made, by us or by creatures like us, is selective. In this section I shall argue that every single causal judgement is selective. And in the next section, I shall explain why someone interested in causation *in the objects* needs to take ubiquitous features of causal *judgements* seriously.⁵

Let us start by saying that a causal judgement is *selective* if it does not mention all the conditions for a given effect. It is, I think, incontrovertible that every singular causal claim is selective. For most of the conditions for any particular effect are beyond the humble scope of our knowledge. Many are in the distant past; of those in the present, many scientific and medical facts remain beyond our collective grasp, or simply beyond our individual education; and there are also many contemporaneous conditions which we could know about but just happen not to — a backup generator which seamlessly preserves you in blissful ignorance of a powercut, for example. There are also conditions which we might know about but which, if asked to produce an exhaustive list, we would simply forget to include; these too are not mentioned in our judgements. And if we agree that absences or omissions can be objects of our causal judgements, then the problems of ignorance and oversight are even clearer, because the number of absences which, had they been presences, would have prevented a given effect is usually unimaginable.

Universal quantification does not count as mention, for these purposes. I do not think the following is a genuine causal judgement:

The causes of my sipping this tea are all those events but for which I would not be sipping this tea.

One way we can tell that this is not really a causal judgement is that it lacks what Schaffer calls the predictive and explanatory signature of causation [Schaffer, 2005, p299]. It does not enable us to make any predictions, other than:

If any one of the events but for which I would not be sipping this tea were not to occur then I would not be sipping this tea.

But that is not a substantive prediction; it is a mere paraphrase of the claim we started with. Nor does the claim offer even a shadow of an explanation as to why I am sipping tea. Consider:

I am sipping tea because of the things but for which I would not be sipping tea.

This is not informative: it does nothing to further our understanding, and therefore it is not a real explanation.

⁵I devote a couple of sentences to this line of argument in [Broadbent, 2008].

The reason that such universally quantified statements do not enable us to make predictions or explanations, and the reason that they thus fail to be causal judgements, is that they are open to the question, “Yes - but what *are* those events, but for which you would not be sipping your tea?” A causal judgement makes a claim about how some particular events are related to some others. The universally quantified statement above does not do that. Rather, it makes a claim about the nature of causation: a claim which is either definitional, and thus irrelevant to our existing notion of causation, or else false, if the argument I am developing is correct.

There are, then, no singular causal judgements which fail to be selective: limits on our knowledge prevent us mentioning them, and universal quantification over all the conditions of a given effect will not yield a causal judgement at all.

What should we say about general causal judgements, such as “Striking matches causes them to light”? I am content to say very little. The counterfactual analysis of causation is concerned with singular causation; it is therefore not clear how to translate the problem of selection, as I have posed it, to the level of general causal claims. If they are quantified claims about singular causation among events, as Lewis hypothesises [Lewis, 1973a, p162], then I am satisfied that all causal judgements remain selective. For I have argued that judgements about singular causation are selective; it is hard to see how quantifying over cases of singular causation will change that, except by the sort of trivialisation which we recently dismissed. When we generalise about the causes of match-lighting, we quantify over other events of the same sort occurring in different instances of causation, not to other events of different sorts related to the self-same particular effect. If a judgement about each of the instances over which we are quantifying is selective, it is hard to see how a judgement about a number of those instances taken together could fail to be selective. For the reason selection is ubiquitous among singular judgements is that our knowledge is limited, and generalising does not remove these limitations.

3 The New Riddle and the Old

The ubiquity of selection in our judgements is enough for me to want to try to explain it. But you might not be similarly impressed. You might be more interested in understanding what it is that makes those judgements true or false. Indeed one of the reasons given by Schaffer, for taking causal selection more seriously than Mill and Lewis do, might move you in exactly the opposite direction: the role of causal selection in moral and legal thinking, and especially the sensitivity of selective judgements to moral and legal contexts, might suggest that causal selection is a feature of our judgements *only*. According to Helen Beebe, “nobody within the tradition of the metaphysics of causation that I’m concerned with here thinks that causal facts depend on human-dependent norms” [Beebe, 2004, p297]. Or you might have some other reason for thinking that explaining causal selection is a different task from explaining causation.

The fact that selection is ubiquitous in causal judgements does not immediately show that an account of what those judgements are about needs also to be an account of selection. But in this section, I shall argue that it *does* show that causal selection needs to be taken seriously, as a problem for causation, and that the radically unselective view of Lewis is guilty of a sort of selective scepticism with regards to the possibility of an objective basis for causal selection. Other ubiquitous features of our causal judgements receive more serious treatment, and I argue that there is no reason to treat selection differently.

Beebe's comment occurs in a discussion of causation by absences or omissions. The discussion focuses on the problem of selection as it arises for causal absences. In effect, Beebe provides a dilemma for those who wish to allow that absences can cause. The problem of selection is particularly vicious for absences in the sense that there are a lot of them; if we allow that the absence of an event can be a condition, in the sense previously defined, of an effect, then there will be a huge number of conditions for every effect. Beebe's question is whether all absences which are conditions are also causes. If we say no, we must provide a way to discriminate between causal absences, and mere condition absences. As I have already hinted, Beebe does not think the prospects for such an account are rosy. She says:

There just *isn't* any objective feature that some absences have and others lack in virtue of which some absences are causes and others are not. So *any* definition of causation by absence that seeks to provide a principled distinction between absences that are and are not causes is bound to fail... [Beebe, 2004, p300]

And the reason, presumably, that we need an objective distinction between cause-absences and condition-absences, is that nobody in Beebe's game thinks that causal facts depend on us humans.

To take the other horn of Beebe's dilemma is to accept that all condition-absences are causes. This, says Beebe, is a wildly counterintuitive view, which does great violence to common sense. Of the failure of a certain Flora to water her orchids, Beebe writes:

The number of possible events, or combinations of events, that are such that, had they occurred, the orchids would not have died, is absolutely enormous... I do not think that most people would happily accept that the failure of each of these events to occur was equally a cause of the orchids' death. [Beebe, 2004, p301]

And the reason that common sense matters is that the motivation for accepting causation by absence is to preserve the literal truth of common-sense causal judgements about absences, such as "Flora's failure to water her orchids caused them to wilt".

Beebe clearly feels the force of the problem of causal selection as it applies to absences. But puzzlingly, the discussion does not mention the problem of selection as it applies to *presences*. Yet Beebe's dilemma would work just as

well for presences; and if we follow Beebee, then by parity of reasoning, we ought to deny that presences cause.

The only difference between presences and absences which is relevant to Beebee's argument is that the number of presences which are conditions for a given effect is probably smaller than the number of absences which are conditions. The number of absences only plays a role in the second horn, where Beebee argues that common sense rebels against admitting such a plethora of causes. She does not go into the reasons for this common sense reaction. If the reason is that many of the absences seem far-fetched or unlikely to occur, then it seems that just the same common sense objection could be put to allowing the presence of oxygen as a cause of this match lighting, in this warm dry room. If there were no oxygen present the match would not light, but the supposition that there is no oxygen in the room is far-fetched. (If it were not, I would install oxygen masks.) If the reason has something to do with the unknowability of all the conditions for an effect, and consequently the impossibility of making a complete causal claim about all the causes of a given effect, then by the argument of the last section, that problem afflicts presences too. And if the reason is neither of these, then I am not sure what it might be. The point is that any view on which the number of absences is a bar to admitting them as causes requires a subsidiary demonstration that there is some threshold, above which the number of conditions for an effect becomes problematic; and further, that the number of absences which are conditions for any given effect is always above the threshold, but the number of presences is always below it.

Without such a demonstration, the argument will extend to causation by presences too. And since the upshot of Beebee's dilemma is that absences do not cause, parity of reasoning would compel us to accept that presences do not cause either. If Beebee's argument is an argument against causation by absence, it is also an argument against causation simpliciter.

The choice between the two horns of Beebee's dilemma is just the choice between a selective and an unselective notion of causation, restricted to the domain of absences. If the problem of selection is a problem for causation by absence, then it is a problem for causation simpliciter. The two horns of Beebee's dilemma arise from my claims (1) and (2). On the one hand, I claim (1) that causal selection is a ubiquitous feature of causal judgement. It is this fact which might tempt us to take the first horn, and find an objective grounding for the principles of selection. But then we encounter (2) the fact that selective judgements have no apparent basis in the objects: there does not appear to be any objective difference, to extend Beebee's expression, between causes and conditions. If we take the second horn of the dilemma, and deny that there is any difference between cause and condition, then this will obviously be no problem; but then we must reconcile ourselves with (1), the ubiquity of selection in our causal judgements.

It is for this reason that I assert (3), that (1) and (2) are in tension. I further claim that the tension is worthy of serious philosophical attention, specifically the attention of people who think that causation is worthy of serious philosophical attention in the first place. This sort of metaphilosophical claim is liable

to provoke irritation, and I too am easily provoked by claims of this sort. But I am not claiming merely that a topic is worthy of philosophical attention; I am alleging some *inconsistency* in the common dismissive attitude to causal selection.

Concerning the idea of necessary connection, David Hume says:

It appears that, in single instances of the operation of bodies, we never can, by our utmost scrutiny, discover any thing but one event following another, without being able to comprehend any force or power by which the cause operates, or any connexion between it and its supposed effect. [Hume, 1748, p113]

Hume argues that there is no apparent basis for our ordinary distinction between cause and mere coincidence — between cause-effect pairs, and just “one event following another”.

This importance of this argument to subsequent thinking about causation is obvious; but *why* is Hume’s argument so important? Why does the conclusion matter? What is the significance of arguing that there is no reason to suppose that there is an objective difference between cause and mere coincidence? The reason, obviously, is that we normally think there is a difference between cause and coincidence; we normally do think that causes make their effects happen, and thus that there is some sort of connection between them. The connection between cause and effect is something we make judgements about incessantly. Every causal judgement seems to presuppose some distinction between a cause-effect pair and a pair of events which just happen to have both occurred. To say that one event causes another just *is* to distinguish that pair of events from a coincidentally conjoined pair. That is why we are so struck when Hume points out that there is no apparent basis for this distinction in the objects.

I am arguing that we should understand the significance of Hume’s sceptical challenge to the notion of necessary connection in the following terms.

- (1’) Necessary connection is ubiquitous in our causal judgements; but
- (2’) Necessary connection is apparently without basis in the objects.

These claims are of a similar form to (1) and (2). That is no artifice: Hume really does argue that a basis for necessary connection is not to be found in the objects; and this conclusion would not be startling if the notion of necessary connection were not ubiquitous in our causal judgements. In the case of (1’) and (2’), there is little difficulty in accepting that:

- (3’) (1’) and (2’) are in tension and generate a deep philosophical problem.

I suggest that it is inconsistent to accept (3’) and deny (3). The problem is just the same: a ubiquitous feature of our causal judgements is without apparent basis in the objects. Amazingly, this is exactly the reason that the problem of

selection is frequently dismissed as an irrelevance — what David Armstrong has called “metaphysical fluff”.⁶ But if that is a basis for dismissing causal selection as a mere whim, a pragmatic overlay, an invidious discrimination among causes, then it is also a reason for dismissing the notion of necessary connection as an instance of the same.

Of course, thorough-going scepticism is a respectable way to respond to Hume’s problem. It is equally a respectable way to respond to the problem of causal selection. Another respectable response is to seek to meet Hume’s challenge, by seeking an analysis of our notion of necessary connection in terms whose basis in the objects is less obscure (Hume’s own constant conjunction analysis may be seen in this light). The same goes for the problem of causal selection; we could seek an analysis of the principles governing our selective judgements which seeks to clarify their objective basis. In neither case need we commit to finding a fully objective basis for the distinction in question; but simply trying to clarify what the distinction might be is a big enough step. My charge against the Mill-Lewis position is that it takes one problem seriously, but not the other. This is objectionable because the nature of the problem — the *reason* each is a problem — is the same. In each case, we have a ubiquitous feature of causal judgement which is apparently without basis in the objects. Why should that be a problem for one feature of causation but not another? The lack of an objective basis for causal selection is not a sound reason for philosophers interested in causation to ignore it; on the contrary, it is a provocation which they cannot honourably overlook.

It is open to deny that the source of the problem is as simple as I have made it sound. Maybe there is some special feature of the notion of necessary connection which makes its presence in judgements and apparent absence in objects philosophically interesting. This response is not very convincing, though, because the same structure can be discerned in a huge class of philosophical problems. Philosophical reflection reveals some class of common sense judgements to be ill-founded and arguably false. Hume identified one way in which causal judgements do not appear to be well-founded: they make a distinction between cause and coincidence. But the distinction between cause and condition is another respect in which causal judgements do not appear to be well-founded. It is open to adopt a sceptical attitude concerning either dubious feature of our common sense judgements. It is not altogether reasonable, however, to adopt a selectively sceptical attitude to one but not the other. As we saw, Beebe argues that to salvage causation by absence for the sake of common sense is a Pyrrhic victory if you then offend common sense with an embarrassment of causes. I am suggesting that it is similarly Pyrrhic to salvage the notion of necessary connection with, say, a counterfactual analysis, while at the same time seeking to amputate another limb of the common sense concept you were seeking to defend.

I have gone too far, in that I do not need to defend the thesis that any analysis of causation must also be an analysis of causal selection (even if that

⁶In conversation.

is what I believe). My claim is merely that the problem of selection is important, deserving some analysis or other; and moreover that it is a problem in the metaphysics of causation, and it is not legitimate to simply stipulate that accounts of the metaphysics of causation can ignore or contradict the selective component of our causal judgements. The problem of causal selection has the same form as Hume's problem for causation, and scepticism about causal selection is as serious as scepticism about the notion of necessary connection, in the sense that every causal judgement has both these puzzling features. It is the common structure of Hume's problem and the problem of selection which suggests the title "New Riddle" for the new problem. Unlike Goodman, I do not claim to have solved the respective old problem; but in the remainder of this essay I shall indicate how I think a solution to both problems might go.

4 The Reverse Counterfactual Solution

Since I have argued that the problem of causal selection is of a kind with Hume's problem, it makes sense to approach both problems in a similar way. My analysis of causal selection will, therefore, be a counterfactual analysis. More precisely, I offer a counterfactual analysis of causation, which makes causation selective.

Schaffer points out that selection is *capricious* [Schaffer, 2005, p313], a word he borrows from Mill. Causal selection is highly context sensitive, and any account must handle this context-sensitivity. In one context, I might say that the match-strike caused the flame, but in another I might mention the presence of oxygen. I already mentioned Hart and Honore's example of a manufacturing process normally conducted in a vacuum. They contend that if the process sprung a leak, the presence of oxygen would after all be counted as the cause of the ensuing fire.

One popular approach to the problem of selection is what I call the contrastive strategy [following Menzies, 2004]. This strategy seeks to employ contrastive theories of causal explanation to account for our selective practices more generally. In causal explanation, we often discriminate between cause and condition; yet explanation is clearly context-sensitive, and we make different discriminations in different contexts. Thus the contrastive strategy appears well-placed to account for the context-sensitivity of causal selection.

For example, Peter Lipton proposes that what makes causal selection work is just the same contrastive mechanism which he says makes contrastive causal explanation work. Thus he writes:

...a cause marks a difference between the situation where the effect occurs and a contrasting situation where it does not. [Lipton, 1992, p136]

For example, we say the strike caused the match to light because we are contrasting the lit match with a situation which is similar in respect of the presence of oxygen, the warmth of the room, the dryness of the match, and so on, but where the match is not lit. The strike is a difference between the actual case

and this contrast case. Similarly, when we say the driver caused the pedestrian's injury, we are contrasting the circumstances of the injury with a situation in which no injury occurs. In that contrast case, the pedestrian is still present, but the driver does not swerve onto the pavement. So the driver's swerving is the difference, and hence the cause.

Why is one contrast chosen rather than another? Why is the actual injury contrasted with a case where the driver doesn't swerve, rather than one where the pedestrian is elsewhere? Lipton is explicit that answering this question is not part of his account – at least, not part of his account of explanation. He says:

My goal... is to show how the choice of contrast helps to determine an explanatory cause, not to show why we choose one contrast rather than another. The latter question is not part of providing a model of explanation, as that task has traditionally been construed. It is no criticism... of my account of contrastive explanation that it does not tell us why we are interested in explaining some contrasts rather than others.

There are various reasons why this attitude is not entirely satisfactory when we transfer it to an account of causal selection. The principal difficulty concerns the choice of contrast. Since that choice plays a determining role in the outcome of a selective judgement, we would expect an account of selection to include an account of the principles governing contrast choice. This is especially pressing outside the explanatory context where it is rare to explicitly specify contrasts [Menzies, 2004]. Yet if a contrastive mechanism is indeed at work, the predictability of selection [Schaffer, 2005] suggests that we manage to agree on them without specifying them. How do we do this? Moreover, the freedom to explain what you want is a peculiar to explanation. Moral thinking, or example, does not allow such freedom.⁷

Those qualms aside, how does the contrastive approach measure up to the challenge identified previously, of reconciling the ubiquity of causal selection in our judgements with its apparent absence from the objects? It clearly gives some principles of causal selection, whose objective basis is not at issue. There is nothing troublesome about the contrastive mechanism proposed; it is perfectly objective, in the sense that whether or not a given candidate cause is absent from the contrast class is a perfectly objective matter. But the silence regarding contrast choice leaves the answer incomplete. Until we know what principles govern choice of contrast, we cannot say whether they have some objective basis or not; and thus we have not answered the new riddle, even if we have come a little closer to answering it.

Peter Menzies identifies another shortcoming of the contrastive strategy. He says it is “unsatisfactory from an explanatory point of view”, because it “unnecessarily duplicates the use of the idea of a cause as something that makes a

⁷I have given this line of argument more attention elsewhere [Broadbent, 2008].

difference” [Menzies, 2004, p150]. On the one hand, we have Lewis’s advertisement for his counterfactual analysis as an account of difference-making:

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. Had it been absent, its effects — some of them, at least, and usually all — would have been absent as well. [Lewis, 1973a, pp160–161]

On the other hand, we have the Mackian notion of a cause as making a difference in a much more restricted way, the difference between the actual case and some contrast case:

A causal statement will be the answer to a causal question, and the question ‘What caused this explosion?’ can be expanded into ‘What made the difference between those times, or those cases, within a certain range, in which no such explosion occurred, and this case in which an explosion did occur?’ [Mackie, 1974, p35]

The contrastive strategy seems content to adopt the first sort of difference-making as an analysis of fundamental unselective judgements, and the second notion as an analysis of causal selection. But Menzies complains:

...it would surely be a surprising fact, requiring elaborate explanation, if our framework for conceptualizing causation used in two different but crucial ways the very same idea of difference-making. It would be much more likely that our conceptual framework was developed on the basis of a single fundamental application of this idea. [Menzies, 2004, p151]

Our conceptual framework is a funny old thing, and perhaps Menzies slightly underestimates the chance that it contains wrinkles of this sort. Nevertheless, he is surely right that a unified and simple explanation of difference-making would be preferable to a two-stage explanation: it would abound more in explanatory virtues.

I now propose a counterfactual which is true of causes and not true of mere conditions. The counterfactual is arrived at by reversing the Lewisian counterfactual, to yield:

The Reverse Counterfactual Necessary Condition on Causation

If c causes e then $\sim E > \sim C$.

Call the counterfactual in this necessary condition the *Reverse Counterfactual*. Clearly the Reverse Counterfactual will often be a backtracking counterfactual — one whose antecedent denotes something temporally later than the consequent. Lewis and others have given backtrackers a bad name, but I view them

more favourably. Hold any doubts for now, to see how this goes; we will then know whether the prize is worth revising our semantics for counterfactuals.

I strike a match in a warm, dry room, and it lights. I assert that if the match hadn't lit, then I wouldn't have struck it. In support, suppose you come into the room to find me holding a lighted match. Now ask what you would have thought, had the match not been lit. Unless you have reason to think otherwise, you would assume I had not struck the match. To dispute this is to question the reliability of matches quite generally, any time they light; a perverse scepticism, belied by their usefulness.

Immediately it will be objected that, if the match hadn't lit, perhaps I would still have struck it, but clumsily, without sufficient force or speed. After all, matches don't always light. I think this is not an objection, but indicates an advantage of the account. For what we must say to get round the apparent objection is that, from an unlit match, we would infer that the match was not struck *well* — not hard enough, not fast enough, or some such. These factors are indeed part of the cause of the flame: they are causally relevant. It is no harm if the Reverse Counterfactual identifies them as such. For when we say the match was not struck well, we are not introducing an endless *ceteris paribus* clause, including all the necessary conditions for the flame, as we shall now see.

The Reverse Counterfactual sharply differentiates the striking of the match from the presence of oxygen. There is no reason to suppose that if the match had not lit, there would have been no oxygen in the room. In this particular case of a match lighting, in this warm, dry, airy room, we do not infer the absence of oxygen from an unlit match. So we cannot use that manner of argument to support the claim that, if the match had not lit, the room would have suddenly evacuated. Nor does this counterfactual have any intuitive plausibility.

The contrast can be brought out sharply in worlds-talk. It is, at least, arguable that the nearest worlds⁸ where there is no flame are worlds where there is no match-strike. Whereas it is not plausible that the nearest worlds where there is no flame are worlds where there is no oxygen.

Perhaps the most obvious question we might ask about this proposal is whether it can account for the *flexibility* of selection — what Schaffer calls caprice, the extreme context-sensitivity of selective judgements. To reflect this flexibility, the Reverse Counterfactual must be true of different events in different contexts. It must explain why we are sometimes willing to say that the presence of oxygen caused the flame rather than the match-strike.

The Reverse Counterfactual, like any other counterfactual, depends both on context. Now the question is: Does it do so in the same way as causal selection? Consider a match-strike suitably similar to this one, but occurring in a different context: a chamber that is ordinarily oxygen-free (perhaps as part of some manufacturing process). On this occasion, oxygen has leaked in, so when the match lights we cite the oxygen as the cause of the flame (cf. Hart and Honore 1985, 10). I suggest that this, too, is reflected by the Reverse Counterfactual.

⁸Here and at other points I make what Lewis calls the Limit Assumption, that there is at least one closest world [Lewis, 1973b, pp19–21]. This is purely for the sake of exposition. The Limit Assumption plays no role in the account.

We select the oxygen in this situation because the nearest worlds where there is no flame are ones where the match is still struck, but the oxygen is absent as usual.⁹

So I am not questioning the context-sensitivity of selection. I am not saying that the strike is always the cause of the flame, and the oxygen is never the cause. My account is supposed to supplement the contrastive approach, providing an account of the principles which govern the choice of contrast when none is supplied. It remains open to us to manipulate the contrast directly, as it were, by explicitly stating a contrast. My account is supposed to explain how we can get away without doing that so often.

I think this approach also helps explain why causation is so useful to us in contexts where selection is clearly a primary concern. After all, if causation is not selective, it is a bit of a mystery why we should try to express selective judgements using that concept. When I blame a reckless driver for the injury she caused, I am not merely picking among causes, distinguishing her action from the innocent but nevertheless strictly causal passage of the pedestrian along that stretch of pavement at that time. I am actually using causal concepts to do the picking. I say that the driver caused the injury and that the pedestrian did not. This may, of course, be a confused expression of a moral conviction. But if causation is fundamentally unselective, it is a wonder that expressions of this confused sort have any meaning at all.

What about cases where we select more than one event as the cause? Even accepting that all causal judgements are selective in the sense that all the conditions for a given effect are never mentioned, still we sometimes say that things have more than one cause. In *joint causation*,¹⁰ the Reverse Counterfactual is false of each of the jointly-causal events. Suppose you and I together lift a table. Assuming we are both similarly motivated (or similarly lazy), and assuming we are both up to the task, then there is no particular reason to say that, if the table had not risen, you would not have lifted: for you might have done, and I might not have. The same goes for me. The Reverse Counterfactual is therefore

⁹It might be objected that, in citing the Reverse Counterfactual as an explanation of our selective practices, I am relying on an implied claim that the Reverse Counterfactual is sufficient for causation (whereas I maintain that it is necessary but not sufficient). I owe this objection to Arif Ahmed. The correct reply, I think, is that the objection artificially restricts the explanans. Spelled out, my claim is not that the Reverse Counterfactual explains causal selection: rather, it is that *the fact that the Reverse Counterfactual is a necessary condition on causation* explains causal selection. In offering my analysis as an explanation, I may indeed be relying on a claim that a certain fact (namely, the truth of my analysis) in some sense suffices for causation to be selective. But that reliance does not amount to (and in fact is incompatible with) suggesting that the Reverse Counterfactual itself is sufficient for causal selection. By way of analogy, the claim that counterfactual reliability is necessary for knowledge might be advanced as an explanation for the fact that we tend not to regard counterfactually insensitive beliefs as knowledge (if indeed that is a fact). Clearly, this would not imply that counterfactual sensitivity of belief is also sufficient for knowledge.

¹⁰Joint causation is to be distinguished from overdetermination. Two bullets entering the President's heart at the same time overdetermine his death. It may be that they also jointly cause it. But if so, that is a special case. The more common kind of joint causation occurs when two events jointly cause a third, without either being sufficient for the effect, and thus without any overdetermination. That is the sort of case which I focus on here.

false of both our efforts; yet we presumably consider ourselves to have together caused the table to lift. We distinguish our efforts from the continued solidity of the ground, which is a mere condition of our accomplishment in this case.¹¹

The obvious solution is to say that if the table hadn't risen, then at least one of us would not have lifted, and thus that $\sim O(e) > \sim (O(c_1) \& \dots O(c_n))$ where c_1, \dots, c_n jointly cause e . Unfortunately, however, if we allow that counterfactuals with conjunctions for consequents may satisfy the necessary condition on causation imposed by the Reverse Counterfactual, we have an immediate problem. For any propositions P and Q , $\sim (P \& Q)$ follows from $\sim P$. So any conjunction will qualify as the cause of any effect, provided the real cause is one of the conjuncts. If the match had not lit, then it would not be the case that I struck the match and you scratched your ear. But my strike causes the flame without any assistance from your scratch.

To block this difficulty we could seek to disallow counterfactuals with conjunctive consequents. But that would rule out the proposed account of joint causation. Therefore I suggest we further stipulate that each conjunct must figure ineliminably. So if c_1, \dots, c_n are joint causes of e then (i) $\sim O(e) > \sim (O(c_1) \& \dots O(c_n))$ and (ii) for every non-empty proper subset $\{c_x, \dots, c_y\}$ of $\{c_1, \dots, c_n\}$, $\sim (\sim O(e) > \sim (O(c_x) \& \dots O(c_y)))$. That is, there must be no proper subset of the candidate joint cause which itself meets the condition initially proposed for qualifying as the joint cause. If there is, then the others are eliminable.

For example, you help me lift a table: we jointly cause it to rise. If the table had not risen, then it is not the case that both you and I would have lifted. However we cannot, with confidence, say that you would not have lifted, nor that I would not have lifted. Thus neither of our lifts qualifies on its own (so neither is eliminable). This appears reasonable and intuitive. Moreover it is supported by the reasoning which led us to conclude that the Reverse Counterfactual failed for joint causes: for that turned on noticing that we could not say that if the table hadn't risen, you wouldn't have lifted, because if the table hadn't risen then you might have, but I might not have, and vice versa. Finally, note that the solidity of the ground fails the condition. In the circumstances we would surely agree it is false that, if the table had not lifted, the ground might have given way. Otherwise the menial task would acquire a new urgency.

For completeness, note that the proposed selective account of causation works just as well for the problem of selection as it arises among causal absences. Suppose I fail to water the plant on my window-sill, and it wilts. Normally we would say that it was my failure, not yours, nor the Queen's, which caused the plant to wilt. The Reverse Counterfactual agrees. If the plant had not wilted, I would have watered it; but it is false that if the plant had not wilted, you would have watered it — after all, you probably don't even know where I live. And it seems at least as far-fetched that Her Majesty would have paid a visit. We can support the Reverse Counterfactual, and undermine the others, by the Inference Test. If, in these circumstances, my plant had not wilted, then (if you somehow came to hear of it) you would infer that I had been watering the plant,

¹¹In the following solution I am indebted to Torben Rees.

and you would neither infer that the Queen, nor that you yourself, had watered it.¹²

5 Conclusion

The Reverse Counterfactual has been defended and developed with respect to particular examples, but I think a more principled line of thinking is available. The Reverse Counterfactual answers Menzies' call for a unified account of the way that causes make the difference to their effects. The underlying thought is that it is implausible to maintain that some event caused another without maintaining that its occurrence *made the difference* between the occurrence of the effect and its failure to occur. That thought is widely accepted,¹³ and on its own it does not presuppose a selective notion of causation: Lewis's account can be seen as specifying that an event c makes the difference to e just in case, if c hadn't happened, e wouldn't have happened. It follows that a lot of events make the difference to any given effect.

I have suggested a different understanding of difference-making, however: that an event makes the difference by *being* the difference between the actual case and *the counterfactual scenario where the effect does not occur*. This view seems, to me, implicit to some degree in various other authors: in Mill's Method of Difference [Mill, 1887], in Mackie's remarks about causes and effects being differences relative to a causal field (Mackie 1974), in Lipton's Difference Condition on contrastive causal explanation [Lipton, 2004], and even in Lewis's account of contrastive causal explanation [Lewis, 1986a]. But if the view I propose is correct — if an event makes the difference by *being* the difference between the case where the effect occurs and the counterfactual scenario where it does not — then the Reverse Counterfactual must, in principle, be true of all causes. For it is just the counterfactual supposition described — supposing what would be the case if the effect hadn't occurred.

The Reverse Counterfactual is a necessary condition on causation, and it makes it hard to see how a sufficient counterfactual condition could be provided. Whenever a cause is also a condition for its effect, it will be true both that $\sim O(e) >\sim O(c)$ and that $\sim O(c) >\sim O(e)$. If counterfactual dependence is sufficient for causation then effects cause their causes. It seems, then, that counterfactual dependence is not sufficient for causation.

Depriving counterfactual analysis of a sufficient component is unfortunate, but it could also be seen as relatively unsurprising. In §1 I emphasised the strangeness, not widely discussed, of Lewis's conviction that counterfactual dependence suffices for causation. That view prevents measuring instruments from being counterfactually reliable, and prevents causes from being counterfactually reliable indicators of their effects. We could characterise reliability another way,

¹²I am indebted to Kit Patrick for first alerting me to the possibility that the Reverse Counterfactual might help with selection among causal absences [Patrick, 2005].

¹³For example: Mill 1887, Mackie 1974, Lewis 1973a, Hart and Honore 1985, Lipton 1993, Menzies 2004, Schaffer 2005.

but given the central role of causation in inference, that would be a pity. Other challenges to the sufficiency of counterfactual dependence for causation have recently received attention. For example, Hall claims that our notion of causation as counterfactual dependence is in tension with another conception of causation, which he calls *production*. When they clash, production takes precedence, meaning that counterfactual dependence is not sufficient for causation [Hall, 2004]. And, of course, a view which makes counterfactual dependence sufficient for causation also makes causation unselective; and that, I have argued, is unsatisfactory.

So we should not lament the loss of a sufficient counterfactual condition for causation. Although there must be more to some c causing some e than counterfactual dependence between c and e , it does not follow that there is more to causation than patterns of counterfactual dependence between events generally. The Reverse Counterfactual is true much more rarely than Lewis's. Perhaps this can give us a general direction for causation. Which of a given pair of events causes the other might then be settled by appeal to this general direction, and the way the events in question fit into it. In that case, there would be something more to causation than the counterfactuals between those two events, but that something more might be the pattern of counterfactual dependence between events more generally.

Yet even if it turns out that there is more to causation than counterfactual dependence, even counting general patterns of dependence not intrinsic to causal pairs, this need not deprive counterfactual analysis of all interest. We need not despairingly conclude that causation is something else entirely, which merely has counterfactual entailments. Knowledge is sometimes thought to have a counterfactual component among components of other sorts: satisfying a counterfactual is a necessary, but not a sufficient, condition to know, on such an analysis (cf. Nozick 1981). Causation might be like that: a hybrid of counterfactual and other components.¹⁴ Given its complexity and the diversity of applications to which it is suited and put, that ought not be a surprise. And just as specifying the counterfactual component of knowledge is worthwhile, so is specifying the counterfactual component of causation.

The other big issue connected with the Reverse Counterfactual is that it is a backtracking counterfactual. I have indicated that I think backtrackers can be true, but that some asymmetry persists: they are true more rarely than foretrackers. Unfortunately I cannot here provide a semantics to support this contention. Four points may make this pill less bitter. First, let me stress that I do not deny that the Reverse Counterfactual often does not sound natural to say. I am claiming, not that the Reverse Counterfactual is natural or intuitive, but that it is *true* of causes but not of conditions, and that we are committed — perhaps unwittingly — to its truth when we make a causal judgement. Second, Lewis's strong claim that counterfactual dependence is asymmetric has

¹⁴Caroline Sartorio's approach in one of her papers leaves such a possibility open: "the view that I defend here is not an analysis of causation. It sets a constraint on the concept of cause, and thus it helps to carve up the concept, while at the same time leaving some room for different ways of pinning it down" (Sartorio 2005, 71).

been subjected to some fierce criticism.¹⁵ It may be, then, that a retreat from that very strong asymmetry is independently motivated. My analysis of selection could be also be seen as a contribution, of a rather preliminary sort, to an investigation into the circumstances in which backtracking appears to be permissible.

Third, the contrastive account of causal selection, and contrastive accounts of causal explanation too, tend to require backtracking. On Lewis's view, contrastive explanations supply "information... about the actualized history of the explanandum and the unactualized causal history of its unactualized alternatives". As Menzies points out [Menzies, 2004, pp149–150], without backtracking, the unactualized causal history of the unactualized match-strike will be just the same as the actual causal history of the actual match strike. There will then be a small miracle by which the match fails to light. Lewis's own account of contrastive causal explanation therefore invokes backtracking. And in general, contrastive accounts of both explanation and selection more generally seem committed to backtracking in at least a limited way, because they ask us to reason back, in a certain contrast case, from the absence of the effect to earlier differences. So my suggestion is not after all such a grand departure in respect of employing backtracking.

Fourth and finally, the Reverse Counterfactual promises to explain more than just causal selection. Since it does not make causes counterfactually necessary for their effects, it offers good prospects for dealing with causal redundancy — preemption, for example. In general, preempted events also satisfy the Reverse Counterfactual: if the president hadn't died, then neither the assassin nor his backup would have fired. This is a different problem from the one facing Lewis: it is a problem of too much truth, rather than of falsity. Too much truth is not such a bad problem for a necessary condition: we simply need to locate further requirements which the preempted events fail. I have elsewhere explained where I think these requirements might be found [Broadbent, 2008].

The advantage of accounting for causal selection with the Reverse Counterfactual as a necessary condition on causation itself, is that this unified approach meets the challenge posed by our new riddle of causation. Causal selection has whatever status causation is supposed to have, on a counterfactual analysis. We still haven't decided exactly what that status is, until we have decided what the objective status of counterfactual dependence is. But we have achieved considerable unification. We have moved closer to reconciling (1) and (2). Causal selection is ubiquitous because the Reverse Counterfactual is a necessary condition on causation, and is selective; and causal selection has whatever objective basis counterfactual dependence has. What that is, I do not propose to decide; but counterfactual dependence is widely considered to be sufficiently objective to provide an analysis of causation in general. We achieve considerable unification if we can put causal selection on a similarly objective footing.¹⁶

¹⁵For example, by Jonathan Bennett, Huw Price and Adam Elga [Bennett, 2001, 2003, Price, 1992a,b, 1996, Elga, 2000].

¹⁶The immediate ancestor of this paper was presented under a different title at the first meeting of the European Philosophy of Science Association in Madrid in 2007. As well as the

References

- Helen Beebe. Causing and nothingness. In J. Collins, N. Hall, and L.A. Paul, editors, *Causation and Counterfactuals*, pages 291–308. MIT Press, Cambridge, Massachusetts, 2004.
- 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.
- Jonathan Bennett. *A Philosophical Guide to Conditionals*. Oxford University Press, Oxford, 2003.
- Alex Broadbent. Reversing the counterfactual analysis of causation. *International Journal of Philosophical Studies*, 15:169–189, 2007.
- Alex Broadbent. The difference between cause and condition. *Proceedings of the Aristotelian Society*, in press, 2008.
- David Coady. Preempting preemption. In Hall N. Collins, J. and L.A. Paul, editors, *Causation and Counterfactuals*, pages 325–340. MIT Press, 2004.
- Adam Elga. Statistical mechanics and the asymmetry of counterfactual dependence. *Philosophy of Science (Proceedings)*, 68:S313–S324, 2000.
- 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, 2004.
- H.L.A. Hart and A. Honore. *Causation in the Law*. Clarendon Press, Oxford, second edition, 1985.
- David Hume. *An Enquiry Concerning Human Understanding*. Clarendon Press, Oxford, 1748. This edition: 1902, ed. L. A. Selby-Bigge.
- David Lewis. Causation. *Journal of Philosophy*, 70:556–567, 1973a. Page numbers refer to (Lewis 1986).
- David Lewis. *Counterfactuals*. Harvard University Press, Cambridge, Massachusetts, 1973b.
- David Lewis. Counterfactual dependence and time’s arrow. *Noûs*, 13:455–476, 1979. Page numbers refer to (Lewis 1986).

helpful members of that audience, I am grateful to Arif Ahmed, Alexander Bird, Katherine Hawley, Dan Heard, Martin Kusch, Peter Lipton, Peter Menzies, Arash Pessian, Martin Peterson, Huw Price, Stathis Psillos, Amit Pundik, Torben Rees, and the audience at the 2007 Joint Session of the Mind Association and the Aristotelian Society, which heard a related paper. I am also grateful for the generous support of the PHG Foundation.

- David Lewis. Causal explanation. In *Philosophical Papers, Volume II*, pages 214–241. Oxford University Press, Oxford, 1986a.
- David Lewis. Events. In *Philosophical Papers, Volume II*, pages 214–269. Oxford University Press, Oxford, 1986b.
- 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.
- Peter Lipton. Causation outside the law. In H. Gross and T.R. Harrison, editors, *Jurisprudence: Cambridge Essays*, pages 127–148. Oxford University Press, Oxford, 1992.
- Peter Lipton. Making a difference. *Philosophica*, 51(1):39–54, 1993.
- Peter Lipton. *Inference to the Best Explanation*. Routledge, London and New York, second edition, 2004.
- John Mackie. *The Cement of the Universe*. Oxford University Press, Oxford, 1974.
- 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.
- John Stuart Mill. *A System of Logic, Ratiocinative and Inductive*. Harper and Brothers, New York, eighth edition, 1887.
- Robert Nozick. *Philosophical Explanations*. Harvard University Press, Cambridge, Massachusetts, 1981.
- Kit Patrick. Causation by absence. Dissertation for University of Cambridge Natural Sciences Tripos Part II., 2005.
- Huw Price. Agency and causal symmetry. *Mind*, 101(403):501–520, 1992a.
- Huw Price. The direction of causation: Ramsey’s ultimate contingency. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1992:253–267, 1992b.
- Huw Price. *Time’s Arrow and Archimedes’ Point*. Oxford University Press, Oxford, 1996.
- Carolina Sartorio. Causes as difference-makers. *Philosophical Studies*, 123:71–96, 2005.
- Jonathan Schaffer. Contrastive causation. *Philosophical Review*, 114(3):297–328, 2005.