

Contextualising Causation Part I

Julian Reiss*
Durham University

Abstract

This is the first instalment of a two-part paper on the counterfactual theory of causation. It is well known that this theory is ridden with counterexamples. Specifically, the following four features of the theory suffer from problems:

- it understands causation as a relation between events;
- counterfactual dependence is understood using a metric of similarity among possible worlds;
- it defines a non-discriminatory concept of causation; and
- it understands causation as transitive.

A number of philosophers have recently proposed that causation is contrastive because making contrasts explicit defuses counterexamples. A contrastive causal claim has the following form: C rather than C* causes E rather than E*, where C* and E* are alternative or contrast events. In this paper, I show that making contrasts explicit does indeed defuse some counterexamples. However, I also argue that the examples discussed in the literature all share a common feature, viz. that the original causal judgement is ambiguous in one way or another. Contrasting does not help with counterexamples that do not have this feature. Part II of this paper then takes up the hard cases.

1. Introduction

It is well known that the counterfactual analysis of causation is flawed. 40 years after David Lewis lamented ‘It remains to be seen whether any regularity analysis can succeed... without piling on the epicycles’ (Lewis 1973[1993]: 194) and developed a counterfactual analysis to replace the then-standard account, there is plentiful evidence that the counterfactual approach is sharing the fate of its rival: it is piling on the epicycles. Every aspect of the original analysis has been shown to be subject to counterexamples, modifications proposed to fix it, but new counterexamples have emerged and new modifications proposed.¹

It would be a mistake, however, to give up on the account completely as yet. Causation as counterfactual dependence is a major concept of causation, one that is highly relevant for ordinary language as well as historical and legal analyses. As such, it is important and deserves our efforts to get it right.

A number of recent contributions have claimed that causation is contrastive. Specifically, they have argued that causation is not a binary relation between a cause and an effect but rather a ternary or quaternary relation between a cause, an effect and a contrast on the side of the cause, the effect or both (see Hitchcock 1993, 1995, 1996a, 1996b; Maslen 2004; Schaffer 2005, 2010, 2013; Northcott 2008). The usual argumentative strategy is the following. Accounts of causation such as the counterfactual or probabilistic account are shown to be subject to certain counterexamples when it is assumed that causation is binary. It is then argued that the counterexamples disappear when causation is understood to be contrastive and concluded that causation is ternary or quaternary.

I will follow this general strategy focusing on the counterfactual account but show that many counterexamples, even some of those that have been considered by contrastivists previously, cannot be dealt with adequately by making causal claims contrastive alone. In addition, an account of ‘admissibility’ of the chosen contrasts has to be given. It will turn out that certain contextual features determine admissibility and that these features ought to be part of the semantics of causal claims. Once causal claims are relativised to a context, explicit contrasting is often unnecessary.

The article is divided into two parts. Part I introduces the counterfactual account and provides the background and motivation for why one should consider causation a relation of arity greater than two. Readers familiar with the literature on causation should skip ahead to part II in which the difficulties of contrastive causation and the case for contextuality are presented.

2. The Counterfactual Account

The counterfactual account aims to analyse causal judgements regarding individual causal relations (as opposed to causal generalisations). That is, it is an account of *actual* causation.² That claims about actual causation and counterfactual claims are closely related is uncontroversial and very intuitive. Ordinarily, we can test whether or not one event caused another by supposing the earlier event away and checking whether the removal makes a difference to the later event. If I suppose that I would have felt fine had it not been for the fish I ate earlier, I can judge that eating of the fish was a cause of my feeling unwell. This test, sometimes called the ‘but-for test’, plays an important role in the law (Hart and Honoré, 1985). Let us suppose the but-for test is sufficient for causation and define:

The simple counterfactual theory of causation. *C* is a cause of *E* if

- *C* and *E* are actual, distinct events, and
- If *C* were not to occur, *E* would not occur.

The devil is always in the details, however. David Lewis, one of the early contributors to the counterfactual theory of causation, developed his account in response to difficulties he perceived in the then-standard regularity theory. Let us focus on two problems first, the *problem of effects* and the *problem of epiphenomena* (Lewis 1973[1993]: 194). The problem of effects refers to the fact that regularity theories cannot determine which of the two constantly conjoined events is the cause and which is the effect unless time-ordering is explicitly built into the analysis (in which case the earlier event is the cause and the later the effect). Building time into the analysis appeared undesirable to Lewis for metaphysical reasons because it precludes a causal analysis of time and makes backwards causation conceptually impossible.

The problem of epiphenomena refers to the fact that a regularity theory, even one that includes time-ordering, cannot distinguish between a direct causal relation and the relation between the two effects of a common cause. If event *C* is constantly conjoined with later event *E*₁ and with another still later event *E*₂, then *E*₁ will be constantly conjoined with *E*₂ and therefore be judged to be a cause according to the theory.

It is not immediately clear that the counterfactual theory helps with these two problems. Suppose the association between infection with the measles virus (the cause) and the appearance of Koplik spots (the effect) is one-to-one. For an infected person, it seems just as true to say, ‘Had John not contracted the measles virus, he would not have developed Koplik spots’ as it seems to say, ‘Had John not developed Koplik spots, he would not have had contracted the measles virus’. Similarly, if the recent drop in atmospheric pressure caused both the low in the barometer reading and the storm, it would not be unusual to say, ‘Had the barometer reading not been low, there would have been no storm’:(cf. Reiss 2012).

In the first case, an effect can come out as a cause and vice versa according to the simple theory; in the second case, one effect of a common cause can come out as the cause of another effect.

Lewis, like the regularity theorist, therefore has to build an asymmetry into the theory in order to distinguish between cause and effect. He does so by proposing a possible worlds semantics for counterfactual claims, according to which 'Had C not been E would not have been' is non-vacuously true if and only if some not- C world where not- E holds is closer to the actual world than any not- C world in which E holds (Lewis 1973[1994]: 197), and stipulating that possible worlds in which when the cause does not obtain, nor does the effect, are generally closer to the actual world than those in which when the cause does not obtain, the effect still does, while it is not the case that worlds in which when the effect does not obtain, nor does the cause, are generally closer to the actual world than worlds in which when the effect does not obtain, the cause still does.³ Thus, by stipulation, 'Had John not contracted the measles virus, he would not have developed Koplik spots' is true, whereas 'Had John not developed Koplik spots, he would not have had contracted the measles virus' is false and similarly for the barometer case.⁴

Suppose for a moment that the simple counterfactual theory of causation used the bi-conditional 'if and only if' instead of the conditional 'if'; that is, it regarded counterfactual dependence (plus distinctness of events) not only as sufficient but also as necessary for causation. That involves a third problem: that of *pre-emption*. One cause pre-empts another if the second cause would have caused the effect if it had not been for the presence of the first. Two assassins set out to kill a desert traveller, one by poisoning his water and the other by drilling a hole in his flask. As it happens, the traveller dies of dehydration; i.e., his death is caused by the assassin who damaged the flask. But there is no counterfactual dependence: had the assassin not drilled the whole, the traveller still would have died. Now, Lewis solved that problem by making the *ancestral of counterfactual dependence* rather than simple counterfactual dependence necessary for causation: if there is a chain of events $C, D_1, D_2, \dots, D_n, E$ such that each subsequent member counterfactually depends on the previous, then C causes E . Thus, while the traveller's death does not depend directly on the second assassin's action, there is a chain of events: the drilling – dehydration – death such that each subsequent member is counterfactually dependent on the previous one.

I will not discuss the necessary condition any further in this paper because cases of pre-emption and other types of redundant causation (where two or more causes compete in their bringing about an effect) constitute particularly recalcitrant counterexamples to the counterfactual analysis, and making causation contrastive does not even appear to help. What is important in the present context is that the move builds transitivity into the concept of cause. That is, if C causes D and D causes E , then, necessarily, C causes E . Most proponents of counterfactual theories of causation have assumed that causation is transitive. As we will see, however, transitivity is problematic in a variety of cases, and making causation contrastive has been supposed to solve them.

My final remark on the simple counterfactual theory of causation before moving on to the contrastive theory is that its conception of cause is egalitarian or *non-discriminatory*. That is, any condition or factor that makes a difference to whether or not the effect-event obtains, no matter how remote or seemingly unimportant comes out as a cause of the event. Thus, it is not only the bad fish that was a cause of my feeling unwell but also my physical constitution, my parents' meeting and the big bang for that matter. Ordinary language may make distinctions among the different causal conditions, but these are a matter of the pragmatics, not of the semantics of causal claims.

In sum, then, the simple counterfactual theory understands causation as a relation with following characteristics:

- the relation holds between *events*;
- it is one of counterfactual dependence, which in turn is evaluated using a metric of *similarity among possible worlds*;
- the concept is *non-discriminatory*; and
- causation is *transitive*.

3. Counterexamples to the Simple Theory, and How Contrasting is Meant to Help

All four characteristics of the simple theory have been subject to much discussion and criticism (see for instance Collins *et al.* 2004). I will here focus on those counterexamples that have either been dealt with explicitly by proponents of contrastive causation or that will prove illuminating from the point of view defended here. The counterexamples are grouped by the characteristic of the simple theory they pertain to.

3.1. EVENT-CAUSATION

The plausibility of any counterfactual theory of causation depends crucially on the understanding of the notion of events involved in the relation. For instance, if one takes events to be too fine grained, many events that are intuitively not causes of an outcome nevertheless appear as such in the theory. Some far-away supernova will have made a difference to the precise manner of my feeling unwell because of its gravitational influence. But surely we will not want to include it as among the causes. Conversely, if one takes events to be too coarse grained, events that intuitively count as causes will not come out as such. To use an example due to Donald Davidson (1980[1993]: 86): the collapse was caused, not by the bolt's giving way, but its giving way so suddenly and unexpectedly.

One question to be considered is whether and how the timing of an event makes a difference to its standing in a causal relation. Consider an example due to Jonathan Bennett. Bennett is here concerned with what he calls 'Lombard's thesis', which states that an event's time is essential to it (Bennett 1987[1993]: 222–3, *emphasis original*):

There was heavy rain in April and electrical storms in the following two months; and in June the lightning took hold and started a forest fire. *If it hadn't been for the heavy rain in April, the forest would have caught fire in May.*

Add Lombard's thesis to that, and you get

If the April rain hadn't occurred the forest fire wouldn't have occurred.

Interpret that in terms of the counterfactual analysis and you get

The April rains caused the forest fire.

That is unacceptable. A good enough theory of events and of causation might give us reason to accept some things that seem intuitively to be false, but no theory should persuade us that delaying a forest's burning for a month (or indeed for a minute) is causing a forest fire.

Here is a different counterexample⁵:

Leonike's theft (Dretske 1977; *cf.* Hitchcock 1996a, Hitchcock 1996b; Hitchcock 2003). Leonike is interested in purchasing a bike or a scooter she has seen in the window of a nearby store. She lacks

the money to buy either, and one night she breaks into the store to steal the bike, the easier vehicle to take. Later that night she is arrested by the police.⁶

Did Leonike's stealing of the bike cause her arrest? Yes and no. On the one hand, it seems that her theft caused the arrest. Had she acquired the bike in a legal way, she would not have been arrested. On the other hand, stealing *the bike* does not seem to have caused her arrest. Had she stolen the scooter instead, she would still have been arrested.

This scenario is a problem for the simple theory if there is one event ('Leonike's stealing of the bike') and yet two conflicting causal judgements. Dretske's response is to make *event allomorphs* rather than events the relata of causation. According to his theory, there is one event – the stealing of the bike – but two event allomorphs – the *stealing* of the bike and the stealing of *the bike* – and the latter enter into causal relations. Lewis' theory of events also has resources to deal with this problem. According to this theory, there are two events in this case: one which is essentially a stealing and only accidentally involves a bike, and one which essentially involves a bike and only accidentally a stealing.

This case is handled simply and beautifully by the contrastive account.

Contrastive Causation. C rather than C^* causes E rather than E^* if

- C and E are actual events,
- C^* is a possible event alternative to C and E^* a possible event alternative to E , and
- E^* counterfactually depends on C^* .⁷

According to the contrastive account, causal relations are not binary but quaternary. In other words, causal relations obtain relative to an explicit set of contrast events. In the case of Leonike's theft, contrasting helps without multiplying our ontology. There are two relevant alternative events. Relative to C_1^* : the alternative event 'Legal acquisition of the bike', her theft is a cause of the arrest; relative to C_2^* , the alternative event 'Theft of the scooter', her actual theft is not a cause of the arrest. Thus, there are two true causal statements to be made about this scenario:

Leonike's stealing the bike rather than acquiring it legally caused her to be arrested rather than remain free.

Leonike's stealing the bike rather than the scooter did not cause her to be arrested rather than remain free.

In this understanding, there is no absolute fact of the matter whether Leonike's stealing the bike caused her arrest. Rather, there are two true contrastive causal claims.

3.2. SIMILARITY AMONG POSSIBLE WORLDS

The simple theory asks us to consider the closest possible world in which the cause is absent and check whether or not the effect is absent too in this world. Unfortunately, it is not always unambiguous what it means for a cause to be absent. Here's a case due to Clark Glymour (Glymour 1986; cf. Hitchcock 1996b):

Uncle Schlomo's Smoking. Glymour's uncle Schlomo smoked two packs of cigarettes a day and eventually contracted lung cancer. For simplicity, let us suppose that he was never exposed to any other carcinogens except those in his own cigarette smoke. Intuitively, it seems correct to say that his smoking caused the lung cancer. But it may well be the case that in the closest possible world in which Schlomo did not smoke two packs of cigarettes a day, he smoked three packs (because of an addictive personality, say).

Under the simple theory, it is false that ‘Schlomo’s smoking caused his lung cancer’ because in the closest possible world he also contracts lung cancer, so there is no counterfactual dependence. Contrasting makes us see that there are in fact two true causal claims:

Schlomo’s smoking two packs of cigarettes rather than none caused him to contract cancer rather than live healthily.

But:

Schlomo’s smoking two packs of cigarettes rather than three did not cause him to contract cancer rather than live healthily.

Here is a similar case (cf. Schaffer, 2005):

The Railroad Switch. A train has to be diverted from an express to a local track. Track worker P. is responsible for changing the switch. P. is aware of a mechanical obstruction in the switch that frequently leads to derailling trains and therefore exercises extra care in changing the setting. The train arrives safely.

Did the track worker’s action cause the safe arrival of the train? Yes and no. On the one hand, it seems that his action didn’t make a difference. Had he not changed the switch, the train would have arrived safely anyway. On the other hand, because of the mechanical obstruction the train was in fact likely to derail. Only because of the extra care the worker put in, it did arrive safely. Therefore, his action made a difference after all.

Under the simple counterfactual account, the worker’s action comes out as a cause. In the closest possible world, he would have failed and the train would have derailed. A possible world in which the action is not successful is closer than any possible world in which the worker is nowhere near the rails.⁸

Again, contrasting elegantly solves this problem. Relative to C_1^* , the alternative event ‘Obstruct switch’, the actual event is a cause of the safe arrival; relative to C_2^* , the alternative event ‘Set switch to express’, the actual event is not a cause of the safe arrival. Correspondingly, there are two true causal statements to be made about this scenario:

Setting the switch to local rather than obstructing the switch caused safe arrival rather than derailling of the train.

Setting the switch to local rather than express did not cause safe arrival rather than derailling of the train.

There are many more counterexamples to Lewis’ similarity metric (e.g., Fine 1975; Elga 2000), but they relate to the asymmetry of causal relations and are not relevant to the contrastivity issue, so I will not discuss them here.

3.3. NON-DISCRIMINATORY CONCEPT

Lewis, 1973[1993] was very clear in what he was looking for (Lewis.: 1973 [1993]):

We sometimes single out one among all the causes of some event and call it ‘the’ cause, as if there were no others. [...] We may select the abnormal or extraordinary causes, or those under human control, or those we deem good or bad, or just those we want to talk about. 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).

Most of the literature has been following him in this (see in particular the essays of Collins *et al.* 2004). The problem with ignoring the principles of invidious discrimination is that one can do so only at the risk of concept of cause that is entirely irrelevant to ordinary language, history and the law. For the concept of cause that is at work in ordinary language, history and the law *is* selective. It would be uninformative at best and often misleading to cite factors such as the big bang as causes of specific events such as Kennedy's death. Below I will discuss cases where we – ordinary language users, historians and lawyers – don't only fail to assert that some event is a cause but explicitly deny it. At any rate, only one category of counterexamples hinges on the selective concept of cause and almost everything I say would go through for a Lewis-style non-discriminatory concept.

If we understand causation selectively, counterexamples are not hard to come by. A person has been smoking two packs of cigarettes every day of his life and eventually develops lung cancer. It is true that the smoking caused lung cancer and that had he not smoked, he would not have contracted the cancer. It is also true that had he not possessed lungs, he would not have contracted lung cancer but we would not say and perhaps even deny that his possession of lungs caused his cancer (Menzies 2004: 143). A doctor might regard someone's ulcer as a cause of her indigestion, while the person preparing meals sees it in eating parsnips (Menzies.: 2004).

To see how contrasting works in these cases, consider the following example.

The Crooked Bridge. An explanatory plaque near the Crooked Bridge in Mostar, Bosnia, reads: 'The Crooked Bridge collapsed on 31st December 1999 during the winter floods but mainly because of damage inflicted during the war (1992–1995)'.

The event 'collapse of the bridge on 31st December 1999' counterfactually depends on both the winter floods and the war damage. But their relevance shifts with the explanatory inquiry. Those reading the plaque are likely to be tourists who are interested in the question, 'Why did the bridge collapse at all?' By contrast, consider inhabitants of Mostar asking, 'Why did the bridge collapse?' immediately after it happened. They will have got used to the bridge's damage and therefore might be interested in the question, 'Why did the bridge collapse when it did, i.e., on 31st December 1999?'

Correspondingly, there are two contrastive causal claims:

Damaging the bridge during the war rather than sparing it caused the bridge to collapse rather than remain intact.

The river's flooding rather than remaining at low level caused the bridge to collapse in winter 1999–2000 rather than at some other time.

3.4. TRANSITIVITY

The last set of counterexamples concerns the idea that causation must be transitive. This appears not always to be the case. A famous example is due to Michael McDermott (McDermott 1995). A man plans to detonate a bomb. Before realising his plan, a dog bites off his right forefinger. The dog bite causes the right-handed assassin to push the button with his left hand. In turn, the button's being pushed causes the bomb to explode. 'Dog bite' causes 'Left-handed button pushing', 'Left-handed button pushing' causes 'Explosion', but it would be false to say that 'Dog bite' causes 'Explosion'.

Here is another case:

The captain and trainee assassin (Northcott 2008; cf. Hitchcock 2003). Captain yells 'fire', trainee fires. Upon hearing the command, victim ducks. The bullet misses him and victim survives unscathed. The captain's yell caused the victim to duck. Ducking saved his life. But the captain's ordering trainee to kill victim did not cause victim to survive.

Contrastivists argue that two different causal judgements are involved here. We first have:

*Captain's yelling his command (C_1) rather than giving no command (C^*_1) caused victim to duck (E_1) rather than stay upright (E^*_1). Victim's ducking (C_2) rather than staying upright (C^*_2) did not cause him to survive (E_2) rather than die (E^*_2). (The latter negative causal claim is true because in the closest possible world in which the captain remains silent, the trainee does not fire; if trainee does not fire, victim's survival is not counterfactually dependent on his ducking.)*

We second have:

*Captain's yelling his command rather than giving a silent command (C^{**}_1) caused victim to duck rather than stay upright. Victim's ducking rather than staying upright caused him to survive rather than die.*

The failure of transitivity obtains because we illicitly shift from one set of contrasts to the other. We take the ducking to be a cause of survival because of the second causal relation where the contrast event is (C^*_1). But we take the first causal relation to obtain relative to the default of giving no command (C_1). Illicit shifts in contrast events result in failures of transitivity (Schaffer 2005). Without such illicit shifts, a generalisation of transitivity called differential transitivity obtains. Relative to the possible world in which no command is given, victim's action does not make a difference to his survival, so there's nothing to be transitive. Relative to the possible world in which a silent command is given, ducking makes a difference to survival and the relation is transitive: it seems correct to say that the captain's *yelling* his command rather than giving a silent command caused victim to survive rather than die.

4. Analysis

In all four cases where contrasting works the intuitive, non-contrastive causal judgement is ambiguous. Did Leonike's stealing the bike cause her arrest? Yes and no. Yes, because she could have bought the bike, and had she bought the bike she would not have been arrested. No, because she could have stolen a scooter instead, and had she done so she would have been arrested nevertheless. Did Schlomo's smoking cause his cancer? Yes and no. Yes, because he could have refrained from smoking altogether and then would not have contracted cancer. No, because he could have smoked even more and then would have contracted cancer anyway. Did the damage inflicted during the civil war cause the collapse of the crooked bridge on 31st December 1999? Yes and no. Yes, because the damage caused the collapse of the bridge. No, because it was the winter floods that made the difference to its collapse when it happened. Did the captain's yelling his command cause victim's survival? Yes and no. Yes, because his *yelling* the command saved victim's life. No, because his *giving a command* presented a threat to victim's life.

Contrasting works by disambiguating initially mixed causal judgements: relative to one set of contrasts one might choose, one causal judgement is true; relative to another set of contrasts, a different causal judgement. We will see in the second part of this article that 'going contrastive' only appears to solve the problem in many cases. The discussion will reveal that the contrastive theory needs some account of which contrast events are admissible choices and which aren't because it is often easy to find some set of contrasts such that a true

claim about counterfactual dependence issues in a false causal judgement. An important assertion I will make is that this selection of contrast events is part of the semantics and not merely part of the pragmatics of causal judgements. That is, I will argue that choice of contrasts determines what causal judgements are true and false and not (just) which causal judgements are appropriate to make. I will further argue that certain contextual features determine the admissibility of contrast events and therefore that context is part of the semantics of causal judgements.

Short Biography

Julian Reiss is a Professor of Philosophy at Durham University. He has a degree in economics and finance from the University of St Gallen and a PhD in philosophy from the London School of Economics. His main research interests are methodologies of the sciences (especially causality and causal inference, models, simulations and thought experiments, and counterfactuals), philosophy of economics and science and values. He is the author of *Error in Economics: Towards a More Evidence-Based Methodology* (2008), *Philosophy of Economics: A Contemporary Introduction* (2013) and some 40 papers in leading philosophy and social science journals and edited collections.

Notes

* Correspondence: Department of Philosophy, Durham University, 50 Old Elvet, Durham DH1 3HN. Email: julian.reiss@durham.ac.uk

¹ It is ironic in fact that Lewis motivated his search for an alternative to the regularity account by pointing to three problems he thought it faces: the ‘problem of effects’, the ‘problem of epiphenomena’ and the ‘problem of pre-emption’ because it is exactly these problems (and more) from which the counterfactual account now suffers. This is not an accident. Unfortunately, there is no space to go into this matter in any detail here.

² Alternative labels include token or singular causation. ‘Actual’ causation is now generally preferred, for example by Pearl (2000: Ch. 10), Woodward (2003: §2.7), Hitchcock and Knobe 2009, Glymour *et al.* 2010; Glymour and Wembley, 2007.

³ Jonathan Schaffer (personal communication) commented on this point that speaking of Lewis as making stipulations about closeness of possible worlds seems unfair. He thinks that Lewis saw himself as being guided by his linguistic intuitions about which counterfactuals were true and that treating these intuitions as reasons to infer that the closeness facts must be such to support the counterfactuals. I would respond that if it is indeed the case that Lewis was guided by his linguistic intuitions, his intuitions were somewhat non-standard. It is far more plausible, say, that if there is a one-to-one correlation between a disease and its symptoms, as has been assumed here, then to remove the symptoms from the world would mean to remove its cause. Thus, if John hadn’t developed Koplik spots, this must have been due to the fact that he didn’t have measles. For what other reason could it have been the case? Similarly, in the barometer case, it is not only not outlandish to exclaim, when looking at a barometer that says ‘low’ and one had planned a barbecue, ‘If only the barometer had been on “high”, we could have had our barbecue’, it is rather far more intuitive than the Lewisian ‘If the barometer had been on “high”, the weather would have been bad anyway’ (because the change from ‘low’ to ‘high’ was brought about by miracle). Ordinary language counterfactuals are often backtracking (Reiss 2012). Lewis’ intuition might have been different but this doesn’t make them true. It is therefore that I say that he *stipulated* a similarity metric that gets causal judgements right.

⁴ To be sure, Lewis 1986 spells out detailed conditions for what he only stipulated in 1973. That analysis is flawed too, however (see for instance Elga 2000), and its details are irrelevant for the present purposes.

⁵ For easy reference, all cases that are discussed as counterexamples from now on will be inserted and named.

⁶ In the Dretske’s scenario, the thief was named Susan and the alternative loot a pair of skis. As this paper was written in the Netherlands and stolen bikes are a perennial Dutch problem, I gave the thief a Dutch name and substituted a scooter for the skis. Skis are not particularly useful in the Netherlands.

⁷ This is Schaffer’s 2005 and Northcott’s 2008 account. Some accounts are ternary, contrasting only the cause or only the effect. I will use this more general form, and it will be clear from the context which side of the contrast is the relevant one.

⁸ This is certainly true under Lewis 1986 semantics: excising the assassin’s action completely from world history requires a bigger miracle than the small nudge that makes him set the railroad switch badly.

Works Cited

- Bennett, Jonathan 1987[1993], 'Event Causation: The Counterfactual Analysis', *Philosophical Perspectives* 1: 367–86.
Reprinted in Sosa and Tooley 1993: 217–33.
- Collins, John, Ned Hall, and L.A. Paul 2004, *Causation and Counterfactuals*, Cambridge (MA): MIT Press.
- Davidson, Donald 1980, 'Causal Relations', in *Essays on Actions and Events*, Oxford: Oxford University Press, 149–62.
Reprinted in Sosa and Tooley 1993: 75–87.
- Dretske, Fred 1977, 'Referring to Events', *Midwest Studies in Philosophy* 2: 90–9.
- Elga, Adam 2000, 'Statistical Mechanics and the Asymmetry of Counterfactual Dependence', *Philosophy of Science* 68 (PSA 2000): 313–24.
- Fine, Kit 1975, 'Review of Lewis, *Counterfactuals*', *Mind* 84: 451–8.
- Glymour, Clark 1986, 'Statistics and Metaphysics', *Journal of the American Statistical Association* 81: 964–6.
- Glymour, Clark, David Danks, Bruce Glymour, Frederick Eberhardt, Joseph Ramsey, Richard Scheines, Peter Spirtes, Choh Man Teng, Jiji Zhang 2010, 'Actual Causation: A Stone Soup Essay', *Synthese* 175: 169–92.
- Glymour, Clark and Frank Wembley 2007, 'Actual Causes and Thought Experiments', in *Campbell*, Joseph Keim, Michael O'Rourke and Harry Silverstein (eds), *Causation and Explanation*, Cambridge (MA): MIT Press.
- Hart, H.L.A. and Anthony Honoré 1985, *Causation in the Law*, 2nd edition, Oxford: Clarendon.
- Hitchcock, Christopher 1993, 'A Generalized Probabilistic Theory of Causal Relevance', *Synthese* 97: 335–64.
———. 1995, 'The Mishap at Reichenbach Fall: Singular vs. General Causation', *Philosophical Studies* 78: 257–91.
———. 1996a, 'The Role of Contrastive Stress in Causal and Explanatory Generalizations', *Synthese* 107: 395–419.
———. 1996b, 'Farewell to Binary Causation', *Canadian Journal of Philosophy* 26(2): 267–82.
———. 2003, 'Of Humean Bondage', *The British Journal for the Philosophy of Science* 54: 1–25.
- Hitchcock, Christopher and Joshua Knobe 2009, 'Cause and Norm', *Journal of Philosophy* CVI(11): 587–612.
- Kincaid, Harold 2012, *The Oxford Handbook of Philosophy of Social Science*, Oxford: Oxford University Press.
- Lewis, David 1973 [1993], 'Causation', *Journal of Philosophy* 70: 556–67. Reprinted in Sosa and Tooley 1993: 193–216.
———. 1986, 'Postscripts to "Causation"', in *Philosophical Papers*, volume II, Oxford: Oxford University Press, 214–40.
- Maslen, Cei 2004, 'Causes, Contrasts, and the Nontransitivity of Causation', in Collins *et al.* 2004: 341–58.
- McDermott, Michael 1995, 'Redundant Causation', *The British Journal for the Philosophy of Science* 46(4): 523–44.
- Menzies, Peter 2004, 'Difference-Making in Context', in Collins *et al.* 2004: 139–80.
- Northcott, Robert 2008, 'Causation and Contrast Class', *Philosophical Studies* 139: 111–23.
- Pearl, Judea 2000, *Models, Reasoning and Inference*, Cambridge: Cambridge University Press.
- Reiss, Julian 2012, 'Counterfactuals', in Kincaid 2012: 154–83.
- Schaffer, Jonathan 2005, 'Contrastive Causation', *Philosophical Review* 114(3): 297–328.
———. 2013, 'Causal Contextualisms', in Blauw, Martijn (ed.), *Contrastivism in Philosophy*, New York: Routledge, 35–63.
———. 2010, 'Contrastive Causation in the Law', *Legal Theory* 16: 259–297.
- Sosa, Ernesto and Michael Tooley (eds) 1993, *Causation*, Oxford: OUP.
- Woodward, James 2003, *Making Things Happen*, Oxford: Oxford University Press.