

Of Miracles and Interventions

Abstract

Lewis (1973, 1979) claims that, for the purposes of delivering a counterfactual analysis of causation, counterfactuals of the form ‘if c hadn’t occurred, then e wouldn’t have occurred’ should be evaluated with respect to those possible worlds in which the non-occurrence of c is realised by a small *miracle* occurring shortly beforehand. Woodward (2003) disagrees. According to him, such counterfactuals ought to be evaluated with respect to those worlds in which c ’s non-occurrence is realised by an *intervention* on whether-or-not c occurs with respect to e . The notion of an intervention, unlike that of a miracle, is causal and so Woodward’s analysis of causation is non-reductive. As one might expect, Woodward claims compensating advantages for his account. In this paper, I argue that these advantages can be had without the appeal to interventions and the consequent sacrifice of potential reductivity.

1. Introduction

According to the tradition originated by Lewis (1973), deterministic causation is to be analysed in terms of counterfactual dependence between distinct, actual events. Lewis (1979) argues that the counterfactuals employed in such an analysis can themselves be analysed in non-causal terms. He is therefore optimistic about the prospects for reduction.

Woodward's (2003, 74-86) theory of token causation is closely related to Lewis's in that he too attempts to elucidate the notion in counterfactual terms. Yet Woodward thinks that the needed counterfactuals must be analysed in terms of the causal notion of an *intervention*. He therefore gives up on reductivity.

Woodward motivates his account by showing that a counterfactual analysis of causation founded upon Lewis's analysis of counterfactuals will lack extensional adequacy (2003, 133-45). Yet, in what follows, I argue that appeal to interventions (or to any other causal notion) is not necessary to remedy the defects of Lewis's analysis. We therefore needn't give up hope of reduction.

The plan is as follows. After reviewing Lewis's analysis of the needed counterfactuals (Section 2), I consider (Section 3) a line of objection to it developed by Elga (2001), Kment (2006), and Wasserman (2006). I then (Section 4) describe an alternative to Lewis's analysis that avoids this objection without appealing to causal notions. After outlining Woodward's rival, interventionist analysis and describing the objections he raises against Lewis (Sections 5 and 6), I show (Section 7) how the alternative developed in response to the Elga-Kment-Wasserman objection can (still without appeal to causal notions) be

extended to deal with Woodward's objections too. The conclusion (Section 8) is that a reductive counterfactual analysis of causation remains on the cards.

2. Lewis's Account

Let c and e be any two distinct, actual events. Then e depends counterfactually and, according to the tradition originated by Lewis (1973), therefore causally upon c just in case the following counterfactual holds:

A. If c hadn't occurred, then e wouldn't have occurred.

On Lewis's (1979) semantics, a counterfactual is true just in case there is a possible world in which both its antecedent and consequent hold that is *overall more similar (closer)* to the actual world than any in which its antecedent holds and its consequent does not (1979, 465; also 1973, 559-60). The vagueness and context-dependence of this relation of comparative overall similarity reflects that of counterfactuals (1979, 465).

Lewis maintains that, for the purposes of delivering a counterfactual analysis of causation, we ought only consider the truth-values that counterfactuals receive in *ordinary* conversational contexts (1979, esp. 459). In ordinary contexts, a counterfactual of form A will normally come out true only for interpretations of c and e such that (in the actual world) e occurs later than c (so that A expresses a *foretracking* counterfactual). By contrast, in certain *special* contexts it may come out true for interpretations of c and e such

that e occurs prior to c , so that it expresses what Lewis (1979, 458; 2004, 78) calls a *backtracking counterfactual*. Taking only the truth values received in ordinary contexts as relevant ensures that counterfactual analyses, which typically take counterfactual dependence to be sufficient for causation, are not in danger of generating spurious cases of backwards-in-time causation.

Lewis claims that the relation of comparative overall similarity that combines with his closest-worlds analysis of counterfactuals to yield the truth-values that counterfactuals receive in ordinary contexts is

“...governed by the following system of weights or priorities.

- (1) It is of the first importance to avoid big, widespread, diverse violations of law.
- (2) It is of the second importance to maximize the spatio-temporal region throughout which perfect match of particular fact prevails.
- (3) It is of the third importance to avoid even small, localized, simple violations of law.
- (4) It is of little or no importance to secure approximate similarity of particular fact” (ibid., 472)

Lewis (ibid., 467-72) argues that, for typical interpretations of c , the similarity measure given by (1)-(4) ensures that the closest non- c worlds are what shall henceforth be called *type 1* worlds. Type 1 worlds match the actual world in history up until just before t_c (the

time at which c occurs in the actual world), at which time a small, localized, simple violation of actual law (a *small miracle*) occurs to prevent c from occurring, but after which they conform to actual law. Since type 1 worlds are consequently not non- e worlds for any e occurring prior to t_c , (except perhaps where e occurs very shortly before t_c) backtrackers come out false. Since type 1 worlds conform to the actual laws of nature from (just before) t_c onwards, so that none of the effects for which c was lawfully necessary occur, many ordinary foretrackers come out true.

Why do type 1 worlds come out as the closest non- c worlds according to the similarity measure given by (1)-(4)? Consider the relevant alternatives. One such alternative – call it a *type 2* world – is where c fails to occur due, not to a small miracle, but to a difference in initial conditions. Since it is assumed that the actual world is deterministic,¹ it follows that a type 2 world differs from the actual world in all of history and contains no spatio-temporal region of exact match of particular fact. Such worlds are more distant than type 1 worlds according to criteria (2) and (3).

Another alternative would be a *type 3* world, which differs from the actual world, including differing in the non-occurrence of c , until just after t_c when a small miracle occurs to secure exact match of particular fact from that point onwards. Whether a type 3 world is closer than a type 1 world will presumably depend upon whether there is more of space-time after than before t_c . Closer than either a type 1 or a type 3 world would be a *type 4* world, in which small miracles occur both just before and just after t_c to secure the

¹ The system of weights (1)-(4) must be complicated somewhat to deliver the right results in the indeterministic case (see Lewis 1986, Postscript D).

non-occurrence of c compatibly with exact match of particular fact throughout the whole of history except for a short interval around t_c .

Crucially Lewis argues (1979, 472-73) that, for ordinary interpretations of c , *there are no type 3 or 4 worlds*. There are, however, type 1 worlds which therefore come out closest. The existence of worlds of type 1, but not of types 3 and 4, is due to an *asymmetry of overdetermination* whereby events are overdetermined (indeed greatly overdetermined) by their future effects, but *not* overdetermined (or at least not greatly overdetermined) by their past causes (ibid., 473-75). In consequence, it is possible to expunge c compatibly with maintaining exact match of particular fact up until just before t_c by means of a small miracle, but only possible to expunge c compatibly with preserving all the future traces of c (and therefore exact match of particular fact from just *after* t_c onwards) by means of a large, widespread and diverse violation of law (a *big miracle*). Such a miracle would make for dissimilarity that outweighs any similarity gain due to securing perfect match of particular fact after t_c .

Consider one final sort of world that might rival type 1 worlds. These are *type 5* worlds, in which the miracle expunging c does not occur *just before* t_c , but rather *at* t_c . The miracle occurring in type 5 worlds just *is* the failure of c to occur in spite of its antecedent determinants. Lewis wants to take type 1 rather than type 5 worlds as relevant to the evaluation of counterfactuals. This is because type 5 worlds involve abrupt discontinuities (1979, 463), whilst the earlier occurrence of the needed miracle in type 1 worlds allows for a more orderly transition. But it is not at all clear that his similarity criteria (1)-(4) deliver this result (see Woodward 2003, 143). After all, type 5 worlds exactly match the actual

world in particular fact over a larger region. Provided no bigger a miracle is required to expunge c compatibly with preserving history up to t_c than to expunge it compatibly with merely preserving history up to *shortly before* t_c , it thus seems that type 5 worlds will come out closest.

On the face of it, this is not a problem for Lewis. Indeed, it may be a good thing. Taking type 5 worlds, rather than type 1 worlds, as closest ensures we get the desirable result that events occurring immediately prior to t_c are counterfactually and therefore causally independent of c (see Woodward 2003, 143; Lewis 1979, 464). In Section 6, some counterbalancing *problems* with taking type 5 worlds as relevant shall be considered. But until then nothing will turn upon the distinction between the two types of world. For simplicity, I will therefore just follow Lewis in taking type 1 worlds as those we'd like to come out as relevant. For present purposes, the more important point is just that worlds of types 2-4 should *not* be taken as relevant. This is important because, if worlds of types 2 or 3 were the relevant ones, backtrackers would come out true. If, on the other hand, worlds of types 3 or 4 were relevant, ordinary foretrackers would come out false.

3. Troubles with Lewis's Account

It has already been observed that the ability of Lewis's semantics plus similarity metric to deliver type 1 worlds as relevant depends crucially upon the (contingent) existence of an asymmetry of overdetermination. It is this asymmetry that ensures the non-existence of worlds of types 3 and 4 (and the existence of worlds of type 1). But Kment (2006) and

Wasserman (2006) have both given examples of cases in which this asymmetry breaks down, whilst Elga (2001) argues that such cases are as widespread as thermodynamically irreversible processes.

The cases to which they point are each of some event c that is not overdetermined by its future effects (so that worlds of types 3 and 4 exist) but which is nevertheless such that the causal (and counterfactual dependence) relations it enters into exhibits the usual temporal asymmetries. The moral is that, even for such interpretations of c , type 1 worlds seem to be the relevant ones to consider, but that Lewis's account fails to deliver this result. Kment (ibid.) suggests that the needed amendment to Lewis's similarity measure would make reference to causal-explanatory facts, threatening to render circular any attempt to analyse causation in terms of counterfactual dependence.

4. A Stipulative Approach

But appeal to causal-explanatory facts is not necessary to deal with Elga-Kment-Wasserman (EKW) style objections. An alternative would be simply to *stipulate* that it is when evaluated with respect to type 1 worlds that counterfactuals of form A are relevant to what causes what. As Lewis (1979, 462-64) observes, to adopt such a stipulative approach would be to abandon the attempt to get the relevance of such worlds to fall out of a general closest-worlds semantics for counterfactuals.² But this lack of generality needn't trouble

² It would also be to abandon the attempt to make room for the in-principle possibility of backwards-in-time causation (Lewis 1979, 464). Yet the counterfactual analyst of

someone simply interested in a way of evaluating counterfactuals that will allow the development of an extensionally adequate, reductive analysis of causation (see Collins et al. 2004, 9).

5. Woodward's Interventionist Account

Woodward (2003) has recently provided a suggestion about how to evaluate the counterfactuals needed by a counterfactual analysis of causation that is quite different both from Lewis's proposal and from the stipulative approach considered in the previous section. He claims that:

“It is a striking feature of the kinds of counterfactuals that are relevant to causal and explanatory claims that ... we require that they be true when their antecedents are realized by interventions” (ibid., 145)

Since Woodward subscribes to a possible-worlds semantics for counterfactuals, his is to be understood as the claim that the truth-value of a counterfactual like A is relevant to whether c is a cause of e only when it is evaluated with respect to (the closest) non- c worlds in which the non-occurrence of c is realised by means of an intervention.

causation may think this a bullet worth biting if the alternatives are circularity or obvious extensional inadequacy.

The notion of an intervention, as Woodward characterises it, has the following three important features. First, it need not be a human activity:

“[A]ny process, whether or not it involves human activities, will qualify as an intervention as long as it has the right causal characteristics.” (ibid., 94)

Second, as this passage also indicates, the notion of an intervention is *causal*.

Consequently, Woodward admits that the counterfactual analysis of causation which he founds upon it is non-reductive (ibid., 22). Third, an intervention that realises the antecedent of a counterfactual like A is (at least where *c* is an actual event) itself non-actual. Consequently, assuming determinism, its counterfactual occurrence is miraculous relative to the laws and initial conditions of the actual world (see ibid., 127-33).

Since an intervention only counts as such in virtue of having the ‘right causal characteristics’, not just any miracle *m* that expunges *c* from a world will count as an intervention. Crucially, to be an intervention upon whether-or-not *c* occurs with respect to *e*, *m* must not exert any *causal* influence upon whether-or-not *e* occurs except by way of its influence upon whether-or-not *c* occurs (ibid., 98).

The notion of an intervention upon whether-or-not one event occurs is thus defined *with respect to another*. So Woodward’s proposal stated more precisely is that the truth of counterfactual A suffices for *c*’s being a cause of *e* (when and) only when A is evaluated with respect to the closest possible world in which *c*’s non-occurrence is realised by an intervention, with respect to *e*, on whether-or-not *c* occurs.

Woodward's arguments for the superiority of his approach to Lewis's involve appealing to certain examples in which Lewis's picks out the wrong worlds as relevant, whilst his picks out the right ones due to the causal constraints on the notion of an intervention (ibid., 139). His examples add to the case established by EKW-style examples for the rejection of Lewis's approach. Yet I shall argue that, like the EKW examples, Woodward's don't establish the need to make reductivity-compromising appeal to causal notions such as that of an intervention. Indeed, I shall argue that the stipulative approach introduced in Section 4 can be developed in such a way as to handle them.

6. Woodward's Objections to Lewis's Account

In arguing for the superiority of his account to Lewis's, Woodward (ibid., 139) requires us to consider counterfactuals of form B:

B. If none of $c_1... c_n$ had occurred, then e wouldn't have occurred.

So far, it has been supposed that the counterfactual analyst of causation need appeal only to counterfactuals of form A, with simple antecedents concerning single events. These can be considered the special case of counterfactuals of form B in which $n = 1$. But Woodward observes that it is plausible that the counterfactual analyst of causation is committed to results concerning counterfactuals of form B generally, and not merely those of the special form A. In general, if a counterfactual of form B comes out true, it seems that standard

counterfactual analyses, on which counterfactual dependence is taken as sufficient for causation, will yield the result that, between them, $c_1 \dots c_n$ caused e (perhaps their mereological sum caused e).

Just as type 1 worlds seem the relevant worlds to consider in evaluating counterfactuals of form A, we might think that the relevant worlds for evaluating counterfactuals of form B are *type 6* worlds, in which each of $c_1 \dots c_n$ is expunged by a small miracle occurring shortly beforehand. Type 1 worlds are the special case of type 6 worlds in which $n = 1$. Taking type 6 worlds as relevant prevents backtracking, and ensures that only ordinary foretrackers come out true, as required by a counterfactual analysis of causation.

Yet Woodward (ibid., 139) observes that there are examples of causal structures for which Lewis's similarity criteria don't seem to pick out type 6 worlds as relevant. One such structure is given in figure 1. Here b , $c_1 \dots c_n$ and e are each distinct, actual events and the arrows are to be understood as representing the causal relations obtaining between them. Thus each of $c_1 \dots c_n$ and e is an independent effect of b . In addition, we are to suppose that b is a *sufficient* cause of each of $c_1 \dots c_n$ and e .




Figure 1

For such an interpretation of $c_1 \dots c_n$ and e , we would like B to come out false. Taking type 6 worlds as relevant appears to yield the desired result. In such worlds, each of $c_1 \dots c_n$ is expunged by a small miracle occurring shortly beforehand (each breaking a causal link from b). Event b still occurs, and e consequently also still occurs.

But where n is large, Lewis's similarity criteria fail to yield type 6 worlds as relevant. For then the n small miracles required to expunge each of $c_1 \dots c_n$ compatibly with preserving b add up to a big miracle (see Lewis 1986, Postscript B). Yet there are worlds in which none of $c_2 \dots c_n$ occur but that don't involve big miracles. In particular, in *type 7* worlds a single, small miracle occurs immediately prior to t_b to expunge b and hence also each of $c_1 \dots c_n$. Although the region of perfect match of particular fact is less extensive in *type 7* than *type 6* worlds, the former don't contain a big miracle. They therefore come out closer. Indeed, for the causal structure in question, it seems *type 7* worlds come out as the *closest* in which none of $c_1 \dots c_n$ occur. Since b doesn't occur in *type 7* worlds, e doesn't occur either and Lewis's analysis therefore yields the undesirable result that B is true (see Woodward op. cit., 140).

Woodward's account, by contrast, gets the case just right. A miracle m that expunges b does not count as an intervention on $c_1 \dots c_n$ with respect to e because it violates the requirement that it must not causally influence e otherwise than by way of its influence on whether-or-not $c_1 \dots c_n$ occur (ibid.). On the other hand, a set of miracles breaking each of the causal links into $c_1 \dots c_n$ qualifies as an intervention because such miracles do not interfere with b . So *type 6* worlds, but not *type 7* worlds, are worlds in which the non-

occurrence of $c_1 \dots c_n$ is realised by an intervention with respect to e . Indeed they are presumably the closest such worlds and are therefore the relevant ones to consider in evaluating counterfactuals like B.

7. Extending the Stipulative Approach

Could we not secure the same result without appealing to the notion of an intervention, by merely generalising the stipulative strategy of Section 4? The idea would be simply to stipulate that the relevant worlds to consider for evaluating counterfactuals of form B are type 6 worlds.

Recall that in Section 3 some doubt was cast on whether it would be desirable to take type 1 worlds, as opposed to type 5 worlds, as relevant to the assessment of simple-antecedent counterfactuals of form A. The question generalises to counterfactuals of form B. Instead of taking type 6 worlds as relevant, perhaps type 8 worlds, in which the non-occurrence of each c_i is realised by a small miracle *at* (rather than *shortly before*) t_{c_i} should be taken as relevant.

Structures like that represented by figure 1 should alert us to a good reason for favouring type 8 over type 6 worlds. For observe that, no matter how strictly we interpret ‘just before’, we cannot be guaranteed that a miracle occurring just before t_{c_i} (as happens in type 6 worlds for each c_i) will not interfere with a common cause of c_i and e (there could be an infinite sequence of common causes of c_i and e , occurring closer and closer to t_{c_i}). So

perhaps we should stipulate type 8 rather than type 6 worlds as relevant. But this leads to problems of its own, as another example of Woodward's shows:

“You are driving on an unfamiliar freeway in the left-hand lane when, unexpectedly, the exit you need to take appears on the right. You are unable to get over in time to exit and as a result are late for your appointment. There are only two lanes, left and right. Driving in the left-hand lane (rather than the right) caused you to be late.” (op. cit., 142)

Consider a counterfactual of form A (equivalently B, where $n = 1$), with c interpreted as your driving in the left lane as you approach the exit, and e as your being late. A type 8 world is one in which, at t_c , your car is in the right-hand lane in spite of the fact that at all times prior to t_c (and after you first came to be in the left lane) it was in the left lane. We might describe such a world as one in which there occurs a miraculous instantaneous shift of your car from the left to the right lane at t_c .

The trouble with taking such a world as relevant (and perhaps the reason Lewis was reluctant to do so) is that the abrupt discontinuity it involves may itself have an impact upon whether e occurs. As Woodward observes:

“[I]t is not at all unlikely that the very occurrence of this miracle will produce effects that will interfere with your exiting. For example, other drivers will be startled and distracted by the sudden appearance of a car in the right-hand lane,

perhaps very close to or in contact with cars that already occupy the right lane. Perhaps a collision will occur or other drivers may swerve or slow down, with the result that your exit is impeded. There will also be a great rush of air into the space in the left lane previously occupied by your car, a similar rush as air is displaced from the right-hand lane, and accompanying loud noises, all of which may also interfere with your exit. So, if this is the relevant world to consider, A may well be false, contrary to the result that we want.” (ibid.; notation modified)

Of course there are worlds (*type 9* worlds) that, like *type 8* worlds, involve c 's non-occurrence due to a miracle at t_c , but in which further miracles prevent the associated abrupt discontinuity from having an independent influence on whether e occurs. In the example, further miracles might prevent the other drivers noticing, your car colliding with others, the noise and air movements, and so on (ibid.).

Lewis isn't entitled to take *type 9* worlds as relevant, for they involve numerous and diverse miracles that presumably add up to a big miracle, and are consequently more distant than *type 8* (or indeed *type 6*) worlds. Woodward claims that, by contrast, his interventionist account implies their relevance (ibid., 144-45). In order to constitute an intervention, a miracle may need to be complex and large, consisting not only of a small miracle at t_c in order to expunge c , but also further miracles to prevent the first miracle having any independent effect upon e . (And, if these further miracles themselves have independent effects upon e , then still further miracles will be required to suppress these.)

In other words, on Woodward's account, the relevant worlds must be taken as those in which a combination of miracles occurs that is such that not only is c eliminated, but also there is no uneliminated independent effect upon e of any of the miracles introduced. Assuming that such a combination of miracles will always be metaphysically possible, it seems Woodward's account will deliver the intuitively correct results.³

Although Lewis's account seems unable to yield type 9 worlds as relevant, I think a version of the stipulative approach can be given that delivers this result without making reference to interventions or to any other causal notion. The benefits of Woodward's approach can therefore be had without the cost of sacrificing potential reductivity.

The version of the stipulative approach that delivers this result is incorporated into the following counterfactual analysis of causation.⁴

(CA) Where c and e are two distinct, actual events, c is a cause of e just in case there is a set of *actual* events or absences, $d_1 \dots d_n$, distinct from c and e and occurring in the temporal interval between them, such that e does not occur in those worlds that (i) match the actual world up until t_c , (ii) involve the occurrence at t_c of the smallest miracle needed to expunge c , and nevertheless (iii) involve the occurrence of each of $d_1 \dots d_n$ (due, for each d_i such that a

³ If this assumption doesn't hold, then it will not be possible to achieve the absence of c by means of an intervention with respect to e , and so the counterfactual A will be vacuously true and there is a risk that Woodward's account will overgenerate.

⁴ The analysis is here restricted to the case in which the putative cause is simple.

miracle is necessary, to the smallest miracle needed at t_{di} to secure its occurrence in spite of the non-occurrence of c).

Where c is your being in the left lane, and e is your being late, there presumably *is* such a set of actual events and absences. This set includes the failure of a collision to occur, the failure of other drivers to swerve or slow down, the non-occurrence of rushes of air into the left lane and out of the right lane, and so on. A world that matches the actual world until you approach the exit, at which time you instantaneously switch from the left to the right lane, and in which all of these events and absences still nevertheless obtain is a type 9 world, and one in which you exit without inhibition.

The requirement that, for the case to be one of causation, there must be a set of actual events or absences meeting the condition described in (CA) is not ad hoc. As shown by Hitchcock (2001) and Yablo (2002, 2004), such a requirement is required to deal with cases of pre-emption. Example: I shoot and kill Victim but, if I hadn't shot, you would have shot and killed Victim. Here there's no counterfactual dependence of Victim's death on my shot, even though the latter is intuitively a cause of the former. But there *is* a set of actual events or absences that meets the condition described in (CA): namely the singleton consisting of your *failure* to shoot. Consider a world that matches the actual world up to the time I'm due to shoot, but in which I miraculously fail to shoot, and yet you also miraculously fail to shoot. In such a world, Victim doesn't die. Consequently we get the desired result that my shot is a cause of Victim's death.

8. Conclusion

The counterfactual analysis of causation requires some or other specification of how the counterfactuals to which her account appeals are to be evaluated. Lewis tries to get such a specification to fall out of a general semantics for counterfactuals. Yet EKW-style examples seem to show that the specification he arrives at cannot ground an extensionally adequate analysis of causation. Woodward offers arguments to the same effect, and concludes that the causal notion of an intervention must be appealed to in order to remedy the defects of Lewis's account. Yet it has here been argued that, with a bit of ingenuity, an adequate (and independently motivated) specification of how to evaluate the needed counterfactuals can be given in non-causal terms. A reductive counterfactual analysis of causation therefore remains a reasonable hope.

References

Collins, John, Ned Hall, and Laurie A. Paul, eds. 2004. *Causation and Counterfactuals*.

Cambridge, MA: MIT Press.

Elga, Adam. 2001. "Statistical Mechanics and the Asymmetry of Counterfactual

Dependence." *Philosophy of Science* 68 (Proceedings): S313-S324.

Hitchcock, Christopher. 2001. "The Intransitivity of Causation Revealed in Equations and

Graphs." *Journal of Philosophy* 98 (6): 273-99.

Kment, Boris. 2006. "Counterfactuals and Explanation." *Mind* 115 (458): 261-310.

Lewis, David. 1973. "Causation." *Journal of Philosophy* 70 (17): 556-67.

— — — 1979. "Counterfactual Dependence and Time's Arrow." *Nous* 13 (4): 455-76.

— — — 1986. "Postscripts to 'Counterfactual Dependence and Time's Arrow'." in Lewis,
David. 1986. *Philosophical Papers*. Vol. II. 52-66. Oxford: Oxford University Press.

Wasserman, Ryan. 2006. "The Future Similarity Objection Revisited." *Synthese* 150 (1):

57-67.

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*.

Oxford: OUP.

Yablo, Stephen. 2002. "De Facto Dependence." *Journal of Philosophy* 99 (3): 130-48.

Yablo, Stephen. 2004. "Advertisement for a Sketch of an Outline of a Prototheory of Causation" In Collins, Hall, and Paul. 2004, 119-37.