A Probabilistic Analysis of Causation

Luke Glynn
(luke.glynn@philosophy.ox.ac.uk)

ABSTRACT

The starting point in the development of probabilistic analyses of causation has usually been the naive intuition that, in some relevant sense, a cause raises the probability of its effect. But, as is well-known, there are cases both of non-probability-raising causation and probability-raising non-causation. Most existing attempts to deal with these cases within a probabilistic framework have brought the resultant analyses into tension with the possibilities of direct non-probability-raising causation, intransitive causation, action-at-a-distance, prevention and causation by absence and omission. I show that an examination of the structure of the problem cases suggests a different way of dealing with them, one which avoids the problems associated with existing probabilistic analyses.

1. Introduction

Our best physical theories seem to indicate that at least some of the fundamental laws of nature are probabilistic. Recognition of this fact has motivated attempts to formulate probabilistic analyses of causation. Among the exponents of such analyses are Good (1961a, b, 1962), Reichenbach (1971), Suppes (1970), Lewis (1986c), Menzies (1989), Eells (1991) and Kvart (2004).

The starting point in the development of these analyses has usually been the naïve intuition that, in some relevant sense, a cause raises the probability of its effect. One straightforward way of cashing this out is in terms of an inequality between conditional probabilities, as follows.[1]

[1] Lewis and Menzies prefer instead to analyse the relevant probability-raising relation in terms of the truth of certain subjunctive conditionals with consequents about unconditional probabilities. I shall not discuss the relative merits of their proposal here.

Take any two distinct, actual events c and e. Suppose that c occurs at time t_c and let $t_{c-\varepsilon}$ be a time immediately prior to t_c. Finally, let capital '*C*' and capital '*E*' be the propositions, respectively, that 'c occurs' and that 'e occurs'. Then a naïve probabilistic analysis says that c is a cause of e iff c raises the probability of e in the sense that inequality (1) holds:

$$P_{t_{c-\varepsilon}}(E|C) > P_{t_{c-\varepsilon}}(E|{\sim}C) \qquad (1)$$

This inequality says that, at $t_{c-\varepsilon}$, the probability of e conditional upon the occurrence of c is greater than the probability of e conditional upon the non-occurrence of c. We may assume that the probabilities in question are to be interpreted as objective chances.

As several philosophers (including Hesslow 1976; Salmon 1980; Edgington 1997; and Hitchcock 2004) have observed, this naïve analysis is inadequate. As usual, counterexamples come from two directions. On the one hand, examples are given of causes that fail to raise the probability of their effects, demonstrating that straightforward probability-raising is not *necessary* for causation. On the other hand, there are examples of probability-raising non-causes, demonstrating that straightforward probability-raising is not *sufficient* for causation either.[2]

2. Causes that Fail to Raise the Probability of Their Effects

The following are two examples of causes that fail to raise the probability of their effects. The first is due to Hesslow (1976, p.291), the second is my own.[3]

1. Jane engages in unprotected sex but takes contraceptive pills which prevent her from becoming pregnant. Studies have shown that contraceptive pills can cause thrombosis. But it is also true that pregnancy can cause thrombosis and that

[2] Two potential objections to the naive analysis which are not addressed in this paper are the following:

(I) *The naive analysis is incompatible with causation in deterministic worlds*: In deterministic worlds every actual event has chance 1, so the conditional probability on the RHS of inequality (1) will be undefined and the inequality will not hold (see Lewis 1986c, pp.178-9).

(II) *The naive analysis gives the wrong results in cases where the direction of causation doesn't correspond to the direction of time*:

(a) *The naive analysis is incompatible with backward causation*: If e occurs prior to c so that e is fixed by $t_{c-\varepsilon}$, then the terms on the LHS and RHS of inequality (1) will both equal 1, so that the inequality will not hold.

(b) *The naive analysis gives the wrong results in cases of simultaneous causation*: If c and e occur simultaneously and c is an *effect* of e, the naive analysis will yield the incorrect result that c caused e (since e is not fixed by $t_{c-\varepsilon}$).

In the present paper, I am concerned merely to show how a probabilistic analysis of causation can deal with cases of non-probability-raising causation and probability-raising non-causation that arise even under the assumption of indeterminism and the assumption that the direction of causation corresponds to the direction of time. Addressing the difficulties that arise when one relaxes these two assumptions is a challenge to be taken up in another paper. Another, related, assumption to be made in *this* paper is that, even as late as $t_{c-\varepsilon}$, the probability of e conditional upon the absence of c remains well-defined.

[3] Other examples are given by Good (1961a, p.318), Rosen (1978, pp.607-8), Salmon (1980, p.65) and Eells (1991, pp.281-2).

contraceptive pills reduce the risk of pregnancy. The probabilities might be such that, overall, Jane's taking birth control pills (call this event 'b') fails to increase the probability of her suffering thrombosis (s). That is, it might be the case that inequality (2) holds:

$$P_{t_{b-\varepsilon}}(S|B) \leq P_{t_{b-\varepsilon}}(S|{\sim}B) \qquad (2)$$

But we have already acknowledged that contraceptive pills *can* cause thrombosis. And suppose that this is such a case: Jane's taking birth control pills causes her to suffer thrombosis. (Our evidence for this causal relationship might be the existence of a complete biochemical process connecting the two events.) We therefore have a case of causation without probability-raising.

2. Billy and Suzy are contemplating whether to cross a rickety bridge over a stream. Billy adopts the following policy: he'll wait and see what Suzy does; if Suzy decides not to cross the bridge, Billy will cross it. If on the other hand Suzy decides to cross the bridge, Billy will toss a coin and cross the bridge just in case the coin lands heads. Billy is heavier than Suzy; there is a moderate chance that the bridge will collapse under Suzy's weight, a high chance that it will collapse under Billy's weight and a very high chance that it will collapse under their combined weight. In fact, Suzy decides to cross the bridge. Billy tosses the coin, the coin lands heads, Billy follows Suzy onto the bridge and the bridge collapses. It seems reasonable to count Suzy's crossing the bridge (a) as a (partial) cause of the bridge's collapsing (d). Nevertheless, because of its negative probabilistic relevance to Billy's crossing (which is a more efficacious potential cause), the probabilities could be such that inequality (3) holds:

$$P_{t_{a-\varepsilon}}(D|A) \leq P_{t_{a-\varepsilon}}(D|{\sim}A) \qquad (3)$$

If so, then although Suzy's crossing the bridge is a cause of the bridge's collapsing, it fails to raise its probability.

Attempts have been made to produce a sophisticated probabilistic analysis that can accommodate non-probability-raising causes. For example Good (1961a, b), Menzies (1989, p.656) and Lewis (1986c, p.179) analyse causation not in terms of probability-raising, but in terms of the ancestral of that relation. Their analyses allow that, where c *does not* raise the probability of e, c may nevertheless be a cause of e provided that there is a sequence of events <c, d_1, …, e> such that each event in this sequence raises the probability of its immediate successor.

There are, however, difficulties with this solution. For one thing, it is not clear that such a sequence will always exist: there might be 'direct' non-probability-raising causation (Salmon, 1980, p.65). In addition, since ancestral relations are transitive, analysing causation in terms of the ancestral of the probability-raising relation has the effect of - to quote Hitchcock (2001a, p.275) - "rendering causation transitive by definition". But there are well-known examples in which causal transitivity appears to

fail.[4] For these reasons, I prefer to seek a different solution to the problem of non-probability-raising causes.

In order to successfully do so, it is important to consider the structure of cases of non-probability-raising causation. The structure of our first example can be represented as follows (cf. Menzies, 1989, pp.645-6; 1996, pp.88-9).

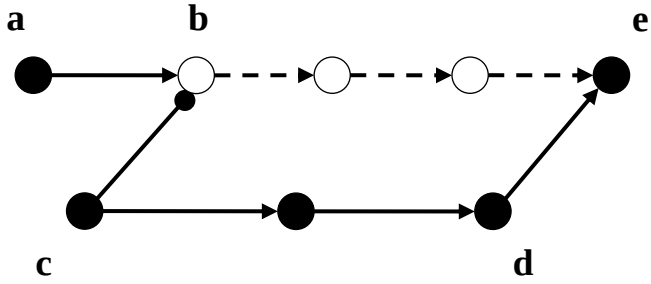

(Figure 1)

In the diagram, solid circles represent neurons that fire, hollow circles represent neurons that fail to fire; solid arrows represent activated stimulatory connections, broken arrows represent non-activated stimulatory connections and the solid arrow ending in a dot represents an activated inhibitory connection (cf. Lewis, 1986c, pp.196, 200-201). It is allowed that both stimulatory and inhibitory connections may work only probabilistically.

The interpretation of the diagram is as follows. Neuron c fires (an event corresponding to Jane's taking the birth control pills) thus initiating a relatively unreliable process which, as luck would have it, runs to completion and causes e to fire (an event corresponding to Jane's suffering thrombosis). The firing of c also prevents the relatively reliable process initiated by the firing of a (which corresponds to Jane's engaging in unprotected sex) from running to completion and causing e to fire. It does so by preventing the firing of b (which corresponds to Jane's becoming pregnant). Intuitively, the firing of c is a cause of the firing of e, but the firing of c fails to raise the probability of the firing of e since its occurrence is negatively relevant to the completion of a more reliable process (the one running from a) by which e might have been brought about.

The structure of the second example is represented in Figure 2:

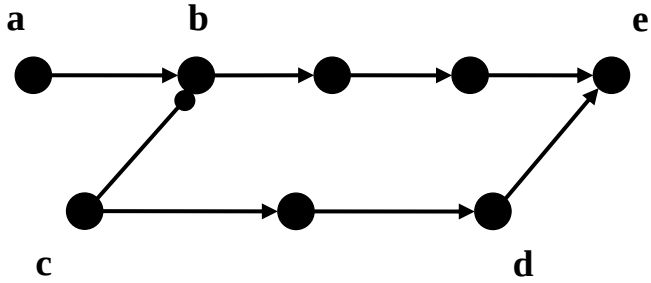

(Figure 2)

[4] Examples are given by McDermott (1995, pp.531-3), Hitchcock (2001a, pp.276-7) and Hall (2004a, pp.183-4; 2004b pp.246-8).

The only difference between this and the previous diagram is that in this case b fires despite the inhibitory connection running from c to b (the inhibitory connection, after all, works only probabilistically and not deterministically). Thus the relatively reliable process (involving Billy's crossing the bridge) runs to completion despite the occurrence of the non-probability-raising cause (Suzy's crossing the bridge). In this case, the firing of b (which corresponds to Billy's crossing the bridge) and the firing of c (which corresponds to Suzy's crossing the bridge) are *joint* causes of the firing of e (the bridge's collapsing).

I have focused on cases of non-probability-raising causation that arise because there are *two* potential processes by which the effect may come about. The diagrams represent the structures of two-process cases. But non-probability-raising causation can also occur in cases where there are *more* than two processes.

What is important is that, in *all* cases of non-probability-raising causation, the non-probability-raising cause c lies on a complete causal process running to the effect e. It is *this* fact that inclines us to say that c is a genuine cause of e. And *given* the fact that c lies on a complete causal process running to e, its negative relevance to alternative processes by which e might have been brought about seems irrelevant to its status as a cause.

Moreover, one can control for the irrelevant inhibitory relationship between c and these alternative processes by means of the following procedure. First, consider those potential alternative processes p_1-p_n to which c is negatively relevant. Some of these processes may be actualised in spite of c, others may not. Nevertheless, for each of these potential processes, p_i, there is at least one possible event b_i that would have lain on this process had the process been fully actualised and which is such that c is *directly negatively relevant* to b_i. To say that c is directly negatively relevant to b_i is to say that c is negatively relevant to b_i and there is no event f that would have lain upon the process p_i which is such that f screens off c from b_i.[5]

Now take the set of those possible events $\{b_1$-$b_n\}$ which lie on potential alternative processes by which e might have come about and to which c is directly negatively relevant. Form a new set (a set of 'facts' or 'conditions' rather than of possible events) which, for every b_i that actually occurs, contains the fact or condition that b_i occurs and, for every b_j that fails to occur, contains the fact or condition that b_j fails to occur. Call this set $\mathbf{s}$, and let capital $\boldsymbol{S}$ be the proposition that all of the conditions in set $\mathbf{s}$ actually obtain. One controls for the irrelevant influence upon the probabilistic relationship between c and e of the negative relevance of c to potential alternative processes by putting $\boldsymbol{S}$ into the condition of our conditional probabilities:

$$P_{t_{c-\varepsilon}}(E|C.\boldsymbol{S}) > P_{t_{c-\varepsilon}}(E|{\sim}C.\boldsymbol{S}) \qquad (4)$$

[5] Definition: f *screens off* c from b_i iff $P_{t_{c-\varepsilon}}(B_i|C.F) = P_{t_{c-\varepsilon}}(B_i|{\sim}C.F)$. The notion of 'screening off' is due to Reichenbach (1971, pp.189-190).

Because non-probability-raising causation arises only in virtue of the negative relevance of the non-probability-raising cause to potential alternative processes, once one controls for this irrelevant influence – by putting ***S*** into the conditions of the conditional probabilities, and so holding fixed whether or not those potential alternative processes were actualised – a relation of positive probabilistic relevance between the cause and effect is revealed.[6]

The procedure just described can be illustrated with reference to the thrombosis example. There is only one potential alternative process (namely, that running from her engaging in unprotected sex) to which Jane's taking the birth control pills is negatively relevant, and only one event upon this process to which Jane's taking the birth control pills is directly negatively relevant: namely, Jane's becoming pregnant. Given that this possible event does not actually occur, take the singleton set containing just the negative condition that Jane fails to become pregnant. Call this set **a**. One controls for the irrelevant inhibitory relationship between Jane's taking the birth control pills and Jane's becoming pregnant by putting the proposition that all of the conditions in **a** obtained into the condition of the relevant probabilities, as in inequality (5):

$$P_{t_{b-\varepsilon}}(S|B.\boldsymbol{A}) > P_{t_{b-\varepsilon}}(S|{\sim}B.\boldsymbol{A}) \qquad (5)$$

Since Jane's taking the birth control pills fails to raise the probability of her suffering thrombosis *only* because of its negative relevance to her becoming pregnant, controlling for this – by putting ***A*** into the condition of the relevant conditional probabilities and thereby holding fixed the fact that she failed to become pregnant – reveals a relation of positive relevance.

In general, because of the common structure of cases of non-probability-raising causation, one can always follow the recipe given above to construct a set of actually-obtaining conditions **s** *relative to which* a non-probability-raising cause c raises the

[6]This method of identifying a positive probabilistic contribution of c to e, by holding fixed events (or failures of events) on alternative processes by which e might have been brought about, is similar to the methods employed in the causal modelling and related philosophical literature (e.g. Pearl, 2000, pp.126-7, 164; Hitchcock, 2001a, pp.284-7; Woodward 2003, pp.45-61, 70-84) for identifying a cause's "*component effect along a causal route*" (Hitchcock, 2001b, p.362, italics original) and for distinguishing this from its 'net' or 'total' effect.

There are differences, however. On my account the positive contribution of a cause to its effect is isolated merely by using conditional probabilities. My method does not appeal to counterfactuals (unlike those of Pearl, Hitchcock and Woodward), to the causal notion of an intervention (Woodward, Pearl), or to a primitive notion of a causal structure or system (Pearl). Nor does my method, like the causal modelling approach, imply that the notion of a cause's contribution along a causal route is relative to the representation of the causal system.

To draw upon the terminology of the causal modelling and related literature one might say that, in this paper, I analyse causation in terms of this notion of a positive 'component effect along a causal route'. It seems to me that, apart from in some limiting cases (in which c only has a positive component effect on e in virtue of neutralizing a threat to e carried along a causal route itself initiated by c - as in *some* prima facie counterexamples to transitivity, e.g. Hitchcock's (2001a) 'Boulder' example), our judgements about whether c was a (token-) cause of e depend unambiguously upon c's having a (positive) component effect upon the probability of e. For further relevant discussion of this issue, see Hitchcock (2001b).

probability of its effect e. Borrowing a term from Kvart (2004, p.360),[7] let us call such a set **s** an *increaser* for c and e.

Since the null set of conditions ∅ (which obtains with probability 1 at all times in all worlds) acts as an increaser in cases of straightforward probability-raising causation, it can be concluded that the existence of an increaser is a necessary condition for causation. The requirement that an increaser exist will replace the requirement of straightforward probability-raising in an adequate probabilistic analysis of causation.[8] [9]

However, given the existence of cases of probability-raising non-causation (in which the null set of conditions acts as an increaser), this is clearly not a *sufficient* condition for causation. It is to the problem of probability-raising non-causation that I now turn.

3. Probability-Raising Non-Causes

Tom and I are playing cricket. Tom hits the ball in the direction of the window. I catch the ball, thus preventing it from hitting the window. Coincidentally, James hits the window with a stone a moment later, breaking it. Tom's hitting the ball raised the probability of the window's breaking and the window did in fact break, but clearly Tom's hitting the ball was not a cause of its breaking.

A feature of this example (and examples of probability-raising non-causation in general[10]) is that the probability-raising non-cause c initiates a causal chain that, if uninterrupted, would (with some probability) have led to the putative effect e. However, the causal chain is cut before it can produce e and e is instead produced in some other way, not involving a causal chain running from c.

Attempts have been made in the literature on probabilistic causation to produce a sophisticated analysis that will exclude probability-raising non-causes from counting as genuine causes. For example, Good (1961a, b) and Menzies (1989, p.656) have

[7] Kvart's definition of this term is somewhat different, as shall be seen in §5 below.

[8] Though to avoid it being trivially true that there exists an increaser for c and e we must exclude from counting as increasers those sets containing any condition x such that the proposition 'x occurred/obtained' is logically related to the proposition 'c occurred' or to the proposition 'e occurred'.

[9] It might be complained that the resultant analysis will fail to be reductive. After all, the recipe given above for constructing an increaser takes the notion of a causal process for granted: in particular, it makes reference to *potential alternative processes* by which e might have been brought about.

My response is as follows: whilst reference to the notion of a causal process is made in giving a recipe for *constructing* (or, rather, *discovering*) an increaser, it is no part of the *definition* of an increaser. That definition is given in purely non-causal, probabilistic, terms. Thus, although the analysis of causation employs the concept of an increaser, it does not thereby import the notion of a causal process.

To put it another way: whether there exists an increaser for c and e is a probabilistic and not a causal fact. But if one wanted to supply someone with a recipe for *discovering* an increaser (not itself an analytic project), then one might as well exploit their existing knowledge of causal processes. Still, if (for some reason) one wanted to supply a recipe that *did not* employ causal terms then one could, with only a little more effort, give one that made reference only to patterns of *probabilistic*, rather than *causal*, dependence.

[10] Other examples are given by Menzies (1989, pp.645-7), Edgington (1997, p.419), and Hitchcock (2004, p.410).

incorporated into their analyses the requirement that there be a *continuous* chain connecting cause to effect. Because cases of probability-raising non-causation involve cut chains, this requirement often succeeds in excluding such cases.

But the requirement of a continuous chain has its costs. One is that it rules out the possibility of action at a spatio-temporal distance *a priori*.[11] A second is that it is far from clear that cases of prevention and causation by absences and omissions involve spatio-temporally continuous chains (cf. Hitchcock, 2004, p.411; Hall, 2004b, pp.243, 249).

So whilst it is true that probability-raising non-causes fail to be genuine causes because the causal chains they initiate are cut, cut causal chains ought not to be characterised in terms of a lack of spatio-temporal continuity. In what follows, I shall attempt to give a more adequate probabilistic characterisation of cut causal chains. The result will be a condition that excludes probability-raising non-causes from counting as genuine causes. This condition supplements the requirement given earlier - that there be an increaser for the putative cause and effect - in such a way that the result is an adequate probabilistic analysis of causation.

3.1. *Cut Causal Chains and Stable Weak Decreasers*

Where a causal chain is cut, it is incomplete. There are certain events (or absences) that would have constituted part of the causal chain had it been complete but which, because the chain is cut, fail to occur.

In the cricket example, my catching the ball cuts the causal chain running from Tom's hitting the ball to the window's breaking. *The ball's hitting the window* is an event that fails to occur but which would have constituted part of the causal chain had it been complete.

The negative conditions corresponding to the failed events on a causal chain running from a potential cause c to a putative effect e have a certain distinctive probabilistic property: they constitute what I shall call 'weak decreasers' for c and e.

Take the negative condition that the ball failed to hit the window ($w^{\sim}$). This acts as a 'weak decreaser' for Tom's hitting the ball (h) and the window's breaking (b) in the sense that, *given* $w^{\sim}$, h fails to raise the probability of b. In other words, inequality (6) holds:

$$P_{t_{h-\varepsilon}}(B|H.W^{\sim}) \leq P_{t_{h-\varepsilon}}(B|{\sim}H.W^{\sim}) \qquad (6)$$

Despite the fact that Tom's hitting the ball raises the probability of the window's breaking, it does not do so when we take the failure of the ball to hit the window as fixed. This is because Tom's hitting the ball is positively relevant to the window's breaking *only because* it is positively relevant to the ball's striking the window.

[11] Indeed, this is one of the reasons for Menzies' later abandonment of his probabilistic analysis (Menzies, 1996, p.94).

Not only is the failure of the ball to hit the window a weak decreaser for Tom's hitting the ball and the window's breaking, it is a *stable* weak decreaser for these two events. That is to say, given the failure of the ball to hit the window, there is no set **t** of actually-obtaining conditions[12] relative to which Tom's hitting the ball is positively relevant to the window's breaking in the sense that inequality (7) holds:

$$P_{t_{h-\varepsilon}}(B|H.W^{\sim}.\boldsymbol{T}) > P_{t_{h-\varepsilon}}(B|{\sim}H.W^{\sim}.\boldsymbol{T}) \qquad (7)$$

If there *were* such a set of conditions, then this would indicate the existence of an alternative causal chain connecting the two events on which the ball's hitting the window does not lie. For example, suppose that John has resolved to hit the window with a hammer just in case Tom hits the ball. Then, relative to the null set of conditions, Tom's hitting the ball raises the probability of the window's breaking even *given* the failure of the ball to hit the window. But this is because of the existence of a second causal chain running from Tom's hitting the ball to the window's breaking (namely that which runs via John's hitting the window). By hypothesis, there is no second causal chain in the original example and the failure of the ball to hit the window therefore acts as a stable weak decreaser.

The existence of a stable weak decreaser for c and e is not a *sufficient* condition of the connecting causal chain's having been cut. This is because *links* on a causal chain running from c to e as well as the *failures* of links on that chain act as stable weak decreasers for the two events. Similar considerations to those advanced above would show that, in our example, the ball's flight as it leaves Tom's bat (which is clearly a *link* and not the *failure* of a link in the causal chain) is a stable weak decreaser for Tom's hitting the ball and the window's breaking.

3.2. *The Distinction Between* Links *and the* Failures *of Links in a Causal Chain*

How, then, are we to distinguish *links* in a causal chain from the *failures* of links in that chain? The answer is that, where d constitutes a *link* in a causal chain running from c to e, there exist *increasers* both for c and d *and* for d and e. On the other hand, where d constitutes the *failure* of a link in a causal chain running from c to e, it will *either* be the case that there is no increaser for c and d *or* that there is no increaser for d and e.

Consider the first claim: where d constitutes a *link* there will exist increasers both for c and d *and* for d and e. This stands to reason. Where d is a link on a causal chain running from c to e, it will be the case both that c causes d and that d causes e. And, as has already been argued, the existence of an increaser for two events is a necessary condition of their being causally related.

This first claim is borne out by the cricket example. The flight of the ball is a *link* in the causal chain. And since Tom's hitting the ball straightforwardly raises the probability

[12] We must exclude sets containing any condition x such that the proposition 'x occurred/obtained' is logically related either to the proposition H or to the proposition B.

of its flight and the flight straightforwardly raises the probability of the window's breaking, the null set of conditions acts as an increaser in both cases.

Now consider the second claim: where d constitutes the *failure* of a link, there will not exist increasers *both* for c and d *and* for d and e. In the example, the ball's failure to hit the window constitutes the *failure* of a link in the chain running from Tom's hitting the ball to the window's breaking. And there is no increaser *either* for Tom's hitting the ball and the ball's failure to hit the window *or* for the ball's failure to hit the window and the window's breaking.

To see that there is no increaser for the ball's failure to hit the window and the window's breaking, consider the following. The ball's failure to hit the window does not raise the probability of the window's breaking. And, as was seen in the earlier discussion of non-probability-raising causes, where one event fails to raise the probability of another, an increaser will exist where there is nevertheless a causal chain connecting the two events. In such cases, the lack of positive relevance is merely a product of the negative relevance of the putative cause to alternative processes by which the effect might have been produced. The increaser is a set of conditions indicating whether or not the alternative processes were actualised.

However, the present case does not have this structure. There is no causal chain connecting the failure of the ball to hit the window to the window's breaking. Hence, no increaser exists. Similar considerations show that there is no increaser for Tom's hitting the ball and the ball's failure to hit the window.[13]

In probabilistic terms, then, the following is the distinction between a *link* and the *failure* of a link in a causal chain running from potential cause c to putative effect e. Suppose that d is an event (or absence) distinct from c and e that is a stable weak decreaser for the two events. Then d is a *link* in the causal chain running from c to e iff there is both an increaser for c and d and an increaser for d and e. On the other hand, d is a *failure* of a link in the causal chain iff there fails to be an increaser for c and d *or* there fails to be an increaser for d and e.

Note that it is not always the case that there is only a single causal chain running from potential cause c to putative effect e. A cause may initiate several chains each of which, if uninterrupted, would (with some probability) bring about the effect. The important question then is how, in probabilistic terms, can one distinguish cases in which *all* of the causal chains running from c to e have been cut (and therefore c fails to be a cause of e)? The answer is that, where they have been, there will be a *set* of events or absences $\{d_1 \ldots d_n\}$ (constituting the failures of links on each of the causal chains running from c to e) that acts as a stable weak decreaser for c and e such that, *for each member - d_i - of the set*,

[13] There are other examples in which there is an event or absence d that constitutes the *failure* of a link in the causal chain running from c to e and in which there exists an increaser for c and d. But, in these cases, there is no increaser for d and e. One such example shall be given in §5. There are also examples in which there is an increaser for d and e but no increaser for c and d.

there is not both an increaser for c and d_i and an increaser for d_i and e. To borrow another term from Kvart (2004, p.366),[14] let us call such a set a *neutralizer* for c and e.

4. Completing the Probabilistic Analysis of Causation

It was noted at the end of §2 that the existence of an increaser for c and e is not a *sufficient* condition for its being the case that c is a cause of e. This is because, as in cases of probability-raising non-causation, it might be the case that each causal chain connecting c and e is cut (in which case there will be a *neutralizer* for c and e). The completed analysis excludes such cases and is as follows.

(C) For any two distinct, actual events (or absences), c and e, c is a cause of e iff:
(i) There is an *increaser* for c and e.
(ii) There is no *neutralizer* for c and e.

Analysis **(C)** overcomes the well-known objections that have been brought to bear against the naïve probabilistic analysis of causation described at the beginning of this paper. In particular, condition (i) is weak enough that **(C)** is satisfied in cases of non-probability-raising causation. Condition (ii) is strong enough to prevent **(C)** from being fulfilled in cases of probability-raising non-causation.

Moreover, unlike existing sophisticated probabilistic analyses,[15] **(C)** overcomes the objections to the naïve analysis in a manner that does not render it incompatible with the possibility of *direct* non-probability-raising causation, intransitive causation, action at a spatio-temporal distance, prevention, or causation by absence and omission. The rest of this section shall be devoted to demonstrating that this is so.

4.1. *Direct Non-Probability-Raising Causation*

The following is a slightly modified version of an example that Salmon gives of 'direct' probability-lowering causation (1980, p.65). Suppose that an unstable atom occupies a state which may be called the 'fourth energy level'. There are several different ways by which it might decay to the 'zeroeth' or ground level. Let $P(m \to n)$ be the probability that an atom in the m^{th} level will make a direct transition to the n^{th} level. And suppose the probabilities are as follows:

$P(4 \to 3) = 0.4$ $\quad$ $P(3 \to 0) = 0.25$
$P(4 \to 2) = 0.4$ $\quad$ $P(2 \to 1) = 0.25$
$P(4 \to 0) = 0.2$ $\quad$ $P(2 \to 0) = 0.75$
$P(3 \to 1) = 0.75$

The probability that the atom will occupy the first energy level conditional upon its occupying the second energy level is 0.25. The probability that the atom will occupy the

[14] Though, as shall be seen in §5, my definition again differs from Kvart's.
[15] With an exception to be noted in §5.

first energy level conditional upon its *not* occupying the second energy level is 0.5.[16] Thus, its occupation of the second energy level lowers the probability of its occupying the first energy level. Nevertheless, it might seem plausible that if the atom goes from the fourth to the second to the first level, then the atom's occupation of the second level is amongst the causes of its occupation of the first level.[17]

This is an example of *direct* probability-lowering causation in the sense that there is apparently no intermediate causal process between the atom's occupation of the second energy level and its occupation of the first energy level. As Salmon says "we cannot, so to speak, 'track' the atom in its transitions from one energy level to another" and, therefore, "it appears that there is no way, even in principle, of filling in intermediate 'links'" of a causal process (Salmon, 1980, p.65). Therefore analyses such as those of Good, Lewis and Menzies, that seek to deal with cases of non-probability-raising causation by replacing the requirement of probability-raising with the requirement that cause and effect stand in the ancestral of that relation will not succeed in yielding the correct result in this case.

The analysis that has been advanced in this paper, by contrast, *does* yield the correct result. This analysis requires neither probability-raising nor the ancestral of probability-raising. It requires merely the existence of an *increaser*. Moreover there *is* an increaser in this case: namely the singleton set containing only the condition that the atom failed to occupy the third energy level. *Given* that the atom failed to occupy the third energy level, its occupation of the second level raised the probability of its occupying the first (since the alternative was for it to decay directly from the fourth to the ground level).

4.2. *Causation by Omission*

Causation by omission is analogous to prevention and to causation by absence both in terms of the difficulties that it poses for existing probabilistic analyses and in the manner in which it is treated by my analysis. The following remarks therefore apply *mutatis mutandis* to these latter two phenomena, which shall not be given separate treatment here.

A farmer's omission to water the crop ($t^{\sim}$) may cause its failure (f). But probabilistic analyses, such as those of Good and Menzies, that require the existence of a continuous causal process connecting cause to effect have difficulty accommodating this fact. For whilst it might be maintained that there *is* an intermediate process connecting $t^{\sim}$ and f – one that involves, for instance, a deficiency of water in the soil, a lack of water being taken up by the roots of the crop, and an insufficiency of water available for metabolic processes – it is doubtful that this process is *continuous*. For instance, it is not at all clear that $t^{\sim}$ has a spatial, or a precise temporal, location. But without such a location, it does

[16] If the atom does not occupy the second energy level, then it has a probability ⅔ of occupying the third level and a probability ⅓ of decaying directly to the ground level (without occupying the first level). The probability of its occupying the first level given its occupation of the third level is 0.75 or ¾. Multiplying ⅔ by ¾, one gets ½ or 0.5.

[17] Salmon (1980, p.65) says that "[a]lthough this example is admittedly fictitious, one finds cases of this general sort in examining the term schemes of actual atoms."

not seem that it can be spatio-temporally contiguous with subsequent events in the process.[18]

The analysis advanced in this paper does not require that cause and effect be connected by a spatio-temporally continuous process. Instead it requires that there be no *neutraliser*. There will exist a neutralizer just in case there is no complete causal chain, but the probabilistic characterization of completeness of a causal chain embodied in the definition of a neutralizer makes no reference to spatio-temporal continuity. Because of this, the analysis can accommodate cases of causation by omission.

Thus in the example there exist stable weak decreasers for t˜ and f, but none of these is a *neutralizer*. The absence of moisture in the soil (m˜) is a stable weak decreaser: *given* that the soil remains dry, t˜ is not positively relevant to f (and is not so relative to any actually-obtaining set of conditions). But t˜ straightforwardly raises the probability of m˜ and m˜ straightforwardly raise the probability of f. The null set of conditions therefore acts as an increaser for m˜ vis-à-vis both t˜ and f, and so m˜ is not a neutralizer for t˜ and f, according to the definition adopted above. The same is true of each of the other links on the causal chain connecting t˜ and f. Indeed, given the completeness of that chain, there is *no* neutralizer for t˜ and f.[19]

4.3. *EPR Phenomena and Action-at-a-Distance*

Because **(C)**, unlike certain other probabilistic analyses of causation, does not require that causes be connected to their effects by continuous processes, it is consistent with an interpretation of EPR phenomena as involving action at a distance.

Suppose that a source emits a pair of entangled electrons that occupy a spin singlet. Electron 1 is sent to destination A, electron 2 to destination B. At A, an observer measures the x-axis spin of electron 1 and finds it to be positive (x_1^+). It is now known with certainty that electron 2 possesses negative spin with respect to the x-axis (x_2^-). One interpretation of this scenario is that the observer's measurement of the x-spin of electron 1 had an instantaneous effect upon the properties possessed by electron 2.

If this interpretation is correct, the case is one of action at a distance. To see that **(C)** is compatible with this interpretation, first observe that the following inequality holds:

$$P_{t_{X_1^+-\varepsilon}}(X_2^-|X_1^+) = 1 > P_{t_{X_1^+-\varepsilon}}(X_2^-|{\sim}X_1^+) \qquad (8)$$

[18] Hall (2004b, p.249) points out that this is not the *only* difficulty for the supposition that causation by omission involves spatio-temporally continuous connecting processes.

[19] The probabilistic analysis (C) is rather liberal when it comes to admitting absences and omissions as causes. In much the same way that it yields the result that the farmer's omission to water the crop was a cause of its failure, it will also yield the result that the absence of rain was a cause and that my omission and your omission to water the crop were causes. I don't regard this liberalness as a defect. It is unproblematic once clear distinctions are made between causation and causal explanation and between causal responsibility and legal and moral responsibility.

The observer's measurement of positive x-spin for electron 1 raises the probability (to 1) of the possession of negative x-spin by electron 2. The null set of conditions therefore acts as an increaser for these two events, so that condition (i) of **(C)** is fulfilled.

Since the influence of the measurement of the x-spin of electron 1 upon the x-spin of electron 2 is not of a sort that involves an intermediate process, there can be neither any *links* nor any *failures* of links in this process. There are therefore no stable weak decreasers and *a fortiori* no neutralizers for x_1^+ and x_2^-. So condition (ii) is also fulfilled and so analysis **(C)** yields the desired result that x_1^+ is a cause of x_2^-.

But what if one doesn't accept that EPR phenomena are correctly interpreted as involving action-at-a-distance? Doesn't analysis **(C)** imply that they do? It does not. Suppose one adopts a hidden variables interpretation instead. For example suppose one holds that, from the time they are emitted, the electrons always possess definite values of x-spin: the correlation between these values is the result of an earlier common cause rather than action-at-a-distance.

Since the common cause occurs no later than $t_{X_1^+ - \varepsilon}$, the properties of the electrons will be probabilistically independent by this time. Therefore inequality (8) will not hold. Indeed this probabilistic independence will be robust in the sense that, not only will the null set of conditions not serve as an increaser for x_1^+ and x_2^-, but there will be *no* increaser for this pair of property-instantiations. Thus, condition (i) of **(C)** is violated. So, if the common-cause interpretation is correct, then **(C)** yields the correct result that the measurement of positive x-spin for electron 1 is *not* a cause electron 2's having negative x-spin.

The analysis advanced in this paper therefore has the desirable property of compatibility with more than one interpretation of EPR phenomena. Whether or not the conditions of the analysis are fulfilled depends upon which interpretation is correct.[20]

4.4. *Failures of Transitivity*

It was noted in §2 that Good, Menzies and Lewis analyse causation in terms of the ancestral of the probability-raising relation and that this has the effect of committing them to the transitivity of causation. It was also noted that there are well-known examples in which causal transitivity seems to fail. One such example is given by McDermott (1995, p.531):

> "*The dog-bite*. My dog bites off my right forefinger. Next day I have occasion to detonate a bomb. I do it the only way I can, by pressing the button with my left forefinger; if the dog-bite had not occurred, I would have pressed the button with my right forefinger. The bomb duly explodes. It seems clear that my pressing the button with my left forefinger was caused by the dog-bite, and that it caused the explosion; yet the dog-bite was not a cause of the explosion."

[20] I have focused upon an example that might be thought to involve action at a spatial distance. But similar considerations show that my analysis also makes room for the possibility of action at a temporal distance.

The probabilistic analysis advanced in this paper, unlike those of Good, Menzies and Lewis, yields the correct result in this case. The dog bite (n) straightforwardly raises the probability of Michael's pressing the button with his left forefinger (l) and l straightforwardly raises the probability of the explosion (x). Hence the null set of conditions acts as an increaser for both event pairs and so condition (i) of **(C)** is fulfilled in each case.[21] Moreover, since the process running from n to l and that running from l to x are complete, there is no neutralizer for either event-pair. The analysis therefore yields the correct result that n was a cause of l and l was a cause of x.

But the analysis *does not* yield the result that n was a cause of x. Event n failed to raise the probability of x, so the null set of conditions does not act as an increaser for this pair. Nor is there a non-null increaser. There would be a non-null increaser if n initiated a causal chain intrinsically conducive to x and if the negative relevance were due merely to the negative relevance of n to one or more alternative processes also conducive to x (as in the examples of §2). If the case had this structure, a set of conditions indicating whether or not these alternative processes were actualised would act as an increaser. But the case does *not* have this structure. In particular it is not as though, on the one hand, n causes l thus initiating a causal process intrinsically conducive to x and, on the other hand, n bears a distinct negative relation to r (Michael's pressing the button with his right forefinger). Rather, it is the same process that both inhibits r and is conducive to l. Insofar as the structure of the case can be represented within our simple neuron diagram framework, it is given in Figure 3:[22]

(Figure 3)

The process connecting n to x runs via i (the incapacitation of Michael's right hand) and l and then on to x. The early stages of this connecting process (such as i) are also the events that pose the threat to x. The negative relevance of n to x is due to n's initiation of this very causal chain. And, consequently, no amount of controlling for events not upon this chain will reveal a relation of positive relevance of n to x. There is therefore no non-null increaser for the event pair and so condition (i) of **(C)** is violated. The analysis thus yields the intuitively correct result that n is *not* a cause of x and that this is a case in which causation fails to be transitive.

[21] Lest it be doubted that l straightforwardly raises the probability of x, recall that it is a stipulation of the example that, when Michael presses the button, his doing so with his left forefinger is the only way he can do it. But suppose it were not; suppose he could have pressed it with some other part of his body (or, for that matter, could have ordered an assistant to press it). Then, assuming that Michael was determined to ensure that the button was pressed one way or another, l may not straightforwardly raise the probability of x. But, in this case, l at least raises the probability of x *given* that Michael didn't press the button with any other part of his body (and *given* that he issued no such order to an assistant). Hence, the set containing only the negative conditions that he didn't do such a thing acts as an increaser for l and x.

[22] See Hitchcock (2001, pp.290-5) for discussion of the structure of this case.

In this section, it has been seen that the probabilistic analysis advanced in this paper gives the correct diagnoses in cases where existing probabilistic analyses run into difficulties. But the cases reviewed here do not only pose problems for existing probabilistic analyses. Many non-probabilistic analyses also deal inadequately with them. For example, Lewis's (1986b) analysis of causation in terms of the ancestral of counterfactual dependence is (like his probabilistic analysis) unable to accommodate intransitivity.[23] On the other hand, process accounts (e.g. Salmon, 1984) and transferred quantity accounts (e.g. Dowe, 2000), have difficulty not only with cases of intransitivity but also with prevention, causation by absence and omission, and action-at-a-distance. This gives us some reason to prefer the analysis presented here to these non-probabilistic rivals (though I shall not, in this paper, attempt a systematic evaluation of the merits of my account vis-a-vis its non-probabilistic rivals).

Not all existing probabilistic analyses run into difficulties in dealing with the types of case considered in this section. There is one sophisticated probabilistic analysis that has recently been advanced which avoids these problems whilst also coming close to adequately treating the problem cases for the naive analysis. It is an account for which I have considerable sympathy and which, in several respects, is close to the analysis that has been advanced in this chapter. For this reason, it merits a more detailed consideration and it is to this task that I now turn, in the final part of the present paper.

5. A Comparison with Kvart's Analysis

Kvart (2004, p.360) recognises that straightforward probability-raising is neither necessary nor sufficient for causation. Like me, he claims that in all cases of non-probability-raising causation there is an 'increaser' for the putative cause c and effect e. However, unlike me, he defines an increaser for c and e as "*an actual intermediate event* [d] that yields probability increase [for c and e] when held fixed" (ibid., my emphasis). Instead of defining an increaser as 'an actual intermediate event', I define it as 'a set of actually-obtaining conditions'. There are good reasons for preferring my definition. These can be seen from a consideration of the following example.

Billy, Suzy and Amy are throwing stones at a bottle. They are all excellent shots. Suzy throws a light stone, Billy and Amy throw heavy stones. The probability of the bottle's breaking conditional upon its being struck by only Billy's stone or only Amy's stone is high. The probability of its breaking conditional upon its being struck by only Suzy's stone is lower, and there is a good chance that Suzy's stone may merely knock the bottle out of the path of Billy and Amy's stones without breaking it. Suzy throws slightly earlier than Billy and Amy, and Suzy's stone hits the bottle, shattering it. Billy and Amy's stones arrive a split-second later at the spot where the bottle had been.

Suppose that the probabilities are such that Suzy's stone's hitting the bottle (h) fails to raise the probability of the bottle's breaking (r). In order for either analysis (Kvart's or mine) to deliver the correct result that h was a cause of r, it must be the case that there is an increaser for h and r. That is, there must be some **b** such that:

[23] Though other counterfactual analyses (e.g. Hitchcock, 2001a) fair better in this regard.

$$P_{t_{h-\varepsilon}}(R|H.\boldsymbol{B}) > P_{t_{h-\varepsilon}}(R|{\sim}H.\boldsymbol{B}) \quad (10)$$

Since I require that **b** be a set of actually-obtaining *conditions*, I can readily allow that there is a **b** such that (10) holds: namely **b** = the set containing just the negative condition that Billy's stone fails to hit the bottle and the negative condition that Amy's stone fails to hit the bottle. *Given that neither Billy nor Amy's stone hits the bottle*, h raises the probability of r. Thus, on my conception of an increaser, there exists an increaser for h and r so that this necessary condition for causation is fulfilled. Kvart, by contrast, has difficulties allowing that there is an increaser in this case. The difficulties arise out of his definition of an increaser as an 'actual, intermediate event'. They are as follows.

First, it is not clear that the failure of Billy's stone to hit the bottle or the failure of Suzy's stone to hit the bottle can be considered *events*. The thesis that there are negative events corresponding to these negative conditions is a controversial one.

Second, both conditions (the failure of Billy's stone to hit the bottle and the failure of Suzy's stone to hit the bottle) are essential components of the increaser that I have identified. Holding fixed merely one or the other is not, on its own, enough to yield a probability-raising relation between h and r. But Kvart's definition implies that an increaser must be a single event rather than a *set* of events. Unless Kvart wishes to allow arbitrary mereological sums of events to themselves count as single events, he will have difficulty finding something that fulfils his definition of an 'increaser'.

Third, it is not at all clear that either the failure of Billy's stone to strike the bottle or the failure of Amy's stone to strike the bottle can be considered *intermediate* between h and r. Indeed, it is in general difficult to assign negative conditions precise locations in space-time.

If there is 'an actual, intermediate event' between h and r that plays the probabilistic role required of an increaser, then Kvart's analysis will deal acceptably with this case. However, so far as I can see there is not. The most obvious candidate for something with the probabilistic properties of an increaser is the one that I have identified. And I think that this is most readily accommodated *as* an increaser if one defines an increaser as a *set of conditions* rather than as a *single intermediate event*.

Kvart, like me, recognises that the existence of an increaser **d** for two events c and e is not a *sufficient* condition for c's being a cause of e. This, according to Kvart, is because there may be "a further decreaser *for it*" (ibid. p.361). That is, there might be an event, f, intermediate between c and e such that:

$$P_{t_{c-\varepsilon}}(E|C.\boldsymbol{D}.F) < P_{t_{c-\varepsilon}}(E|{\sim}C.\boldsymbol{D}.F) \quad (11)$$

Kvart claims that "[f]or [c] to be a cause of [e], it must have an increaser [d] *without* a further decreaser for it" (ibid. p.362). He calls such an increaser *stable*.

In requiring that there be a 'stable increaser' for c and e, I think that Kvart's analysis is equivalent to my own. Note first that, unless it is possible to have an infinite sequence such that for every increaser there is a further decreaser and for every decreaser there is a further increaser (in which case it seems totally unclear what one should say about whether a causal relation exists), the requirement that there be a stable increaser for c and e is equivalent to the requirement that there be no stable decreaser for c and e.

Now condition (ii) of my analysis requires that there be no neutraliser. And a neutraliser (on my definition) is a stable *weak* decreaser, d, for c and e such that there is no increaser for c and d *or* there is no increaser for d and e. A stable weak decreaser for c and e is something relative to which c fails to raise the probability of e no matter what further conditions are added. A stable decreaser is, therefore, one type of stable *weak* decreaser. It is also, I suggest, a type of stable weak decreaser for which, when there exists an increaser for c and e, there never are increasers for it vis-à-vis both c and e.

Since a stable decreaser is a type of stable *weak* decreaser and since (I claim) it is a type for which there cannot be increasers for it vis-à-vis both c and e (when there is an increaser for c and e), in requiring that there be no stable *weak* decreasers that lack increasers vis-à-vis both c and e, I am in effect requiring that there be no stable decreasers for c and e (though clearly this is not all I am requiring). And this requirement is (given a modest assumption) equivalent to Kvart's requirement that there be a stable increaser for c and e.

According to Kvart, even the existence of a stable increaser (and therefore the absence of a stable decreaser) for c and e is not sufficient for its being the case that c is a cause of e. Despite the existence of a stable increaser, c will not be a cause of e if there is a *neutralizer* for c and e (ibid., p.365). A neutralizer, on Kvart's definition, is an event d intermediate between c and e that acts as a stable screener (stable in the sense that there is no other intermediate event that, when conditioned upon, undoes its screening-off of c from e) and that is not *caused* by c (ibid., p.367-8).

The purpose of Kvart's requirement that there be no neutralizer for c and e is, like the corresponding condition of my analysis, an attempt to ensure that the causal chain between c and e is not cut. Kvart holds that an event that constituted a failed link in the causal chain running from c to e would be a stable screener for c and e. However, he recognises that *links* in the causal chain will often also be stable screeners (ibid. p.368n.). It is in an attempt to distinguish *failures* of links from *links* that Kvart introduces into his definition of a neutralizer the requirement that c must not be a *cause* of the putative neutralizer d. A link, d, in a causal chain running from c to e will be such that c causes d. Kvart maintains that, by contrast, the *failure* of a link in a causal chain will not be caused by c.

Whilst Kvart defines a neutralizer as ***an intermediate event*** *that acts as a* ***stable screener*** *for c and e and that is* ***not caused*** *by c*, I define it as ***a set of conditions*** *that acts as a* ***stable weak decreaser*** *for c and e such that for each* d_i *in this set,* ***there do not exist increasers*** *both for c and* d_i *and for* d_i *and e* (the existence of an increaser is taken

to be a necessary condition for causation). Kvart's definition of a neutraliser makes his requirement that there be no neutralizer for c and e more problematic than the corresponding requirement in my analysis. It does so for five reasons.

First, it seems clear that not only events, but also absences, can constitute *failures* of links in a causal chain from c to e. For example, Jane's failure to become pregnant constitutes a failure of a link in the causal chain running from her engaging in unprotected sex to her suffering thrombosis. Because of this, it would be better to define a neutralizers in terms of *conditions* (which can be positive or negative) rather than *events* (since the notion of a negative event is controversial). Moreover, since assigning temporal locations to absences is often problematic, it seems that it would be better to drop Kvart's requirement that a neutralizer be *intermediate* between c and e.

Second, where there are several causal chains running from the potential cause to the putative effect and all of them are cut there will be no single natural event or absence that acts as a stable screener (or as a stable weak decreaser) for c and e. Rather, it will only be sets of events or absences that act as such. For this reasons, it is preferable to define a neutraliser (as I do) as a *set* of events or absences rather than as a *single* event or absence.

Third, on Kvart's definition, a neutralizer is a *stable screener* for c and e, I define it as a *stable weak decreaser*. And because a set of events or absences constituting the *failure* of a link on each causal chain running from c to e will not always be a stable screener, but will always be a stable weak decreaser for c and e, my condition does a better job than Kvart's in serving its purpose: to exclude cases in which all causal chains from c to e are cut.

To illustrate, consider a modified version of the thrombosis example. Suppose that there is a failure of some event a in the causal process running from Jane's taking the birth control pills (b) to her suffering thrombosis (s) (a failure that does not affect the birth control pills' prevention of pregnancy). Call the corresponding absence '$a^{\sim}$'. Even though $a^{\sim}$ constitutes the failure of a link on the only causal chain running from b to s, it is not a screener for b and s, let alone a *stable* screener. *Given* $a^{\sim}$, b is *negatively relevant* and not *irrelevant* to s (because b is negatively relevant to Jane's becoming pregnant). Thus, on Kvart's definition, $a^{\sim}$ does not count as a neutraliser and so Kvart's requirement that there be no neutralizer won't, as it ought to, exclude b from counting as a cause of s.

Nor will his requirement that there must be a stable increaser help here. Jane's failure to become pregnant ($p^{\sim}$) constitutes a stable increaser on Kvart's definition. That is, it is an increaser for which there is no further decreaser. Note that $a^{\sim}$ does not constitute a further decreaser because, given both $p^{\sim}$ and $a^{\sim}$, b is not negatively relevant to s, but *irrelevant* to s. So it seems that Kvart's analysis gives the incorrect result that b is a cause of s.

By contrast, given *my* definition of a neutralizer, condition (ii) of **(C)** *does* exclude b from counting as a cause of s. Though it is not a stable screener, $a^{\sim}$ is a stable weak decreaser for b and s. Moreover, since there is no increaser for b and $a^{\sim}$ or indeed for $a^{\sim}$

and s, condition (ii) of **(C)** is violated when b is taken as a putative cause and s as its putative effect.

Fourth, consider the following version of the 'thrombosis' example. Jane takes birth control pills (b), fails to suffer thrombosis ($t^{\sim}$) and yet dies (d) (from another cause). In this case, $t^{\sim}$ constitutes the failure of a link in the causal chain running from b to d (and is a stable weak decreaser for b and d). Nevertheless, it is plausible that b *caused* $t^{\sim}$ (since b prevented Jane from becoming pregnant and pregnancy is a relatively potent cause of thrombosis). Thus, on Kvart's definition, $t^{\sim}$ does not count as a neutralizer for b and d (even if it is a stable screener). Unless there is some other event or absence that acts as a stable weak decreaser for b and d and which is not caused by b (and there is certainly none that is salient in the presentation of the example), then there is *no* neutralizer for b and d on Kvart's definition and we get the incorrect result that b was a cause of d (since Jane's failure to become pregnant, $p^{\sim}$, acts as a stable increaser for b and d).

By contrast my analysis again yields the correct result. It does so because my definition of a neutralizer doesn't necessarily put *any* constraints upon the relationship between potential cause b and a putative neutralizer $t^{\sim}$. Rather than requiring that it not be the case that b caused $t^{\sim}$, it requires merely that *either* there be no increaser for b and $t^{\sim}$ *or* there be no increaser for $t^{\sim}$ and d. And observe that, whilst the null set of conditions acts as an increaser for b and $t^{\sim}$, there is no increaser for $t^{\sim}$ and d (since there is no causal chain connecting the two events).

Fifth, in contrast to my definition of a neutralizer, Kvart's makes reference to causation: a neutraliser for c and e is a stable screener that is not *caused* by c. As Kvart himself admits (ibid., p.368) this appears, at least *prima facie*, to render his analysis of causation circular.[24] By contrast, given my definition of a neutralizer, my analysis does not even have the appearance of circularity.

Kvart argues that the appearance of circularity in *his* account is illusory (ibid. p.368-9). The idea is that his analysis is to be applied recursively. Suppose that there is a stable increaser for c and e. Then c is a cause of e unless there is a neutralizer for c and e. Suppose that there is a stable screener, d_1, for c and e. Then, unless c is a cause of d_1, d_1 is a neutralizer (on Kvart's definition) and c is not a cause of e. If there is no stable increaser for c and d_1, c is not a cause of d_1 and so not a cause of e. If, on the other hand there *is* a stable increaser for c and d_1, then c is a cause of d_1 unless there is a neutralizer for c and d_1. Suppose that there is a stable screener, d_2, for c and d_1. Then, unless c is a cause of d_2, d_2 is a neutralizer and c is not a cause of d_1 and so not a cause of e. If there is no stable increaser for c and d_2, then c is not a cause of d_2 and so not a cause of e. If, on the other hand, there *is* a stable increaser for c and d_2, then c is a cause of d_2 unless there is a neutralizer for c and d_2. Etc.

[24] There is a similar appearance of circularity in the probabilistic analysis of causation advanced by Eells (1991). By contrast to Kvart, Eells accepts that the circularity in his account is genuine (ibid. pp.350-1), but suggests that it is not vicious (1991, p.351n.) and is apparently content to settle for an analysis that (by contrast to the one advanced here) is not reductive in aspiration.

Ultimately, Kvart observes (ibid., p.369), this sequence might terminate with a stable screener, d_n, for c and d_{n-1} such that there is no increaser for c and d_n, *or* the sequence might not terminate, in which case there will be an infinite regress. If the sequence terminates, then d_n is a neutralizer and so c is not a cause of e. If, on the other hand, no such sequence terminates then since the case is one of prima facie causal relevance of c to e (because there is a stable increaser for c and e) and there is no neutraliser "to overrule the prima facie causal relevance" (ibid.), it follows that c is a cause of e. Because "the infinite regress case is a case of causal relevance" (ibid.), the infinite regress is not vicious.

But even if Kvart is right, the infinite regress is nevertheless an unattractive feature of his account, as is the appearance of circularity. Moreover, it can be avoided by recognising the following. First, any link d_i on the causal chain running from c and e will be such that there is an increaser for c and d_i and an increaser for d_i and e. Second, there will only fail to be *stable* increasers for *both* c and d_i *and* for d_i and e if there is a *failure* of a link in the causal chain running from c to d_i or in that running from d_i to e. Third, any *failure*, d_j, in the causal chain running from c to d_i or in the causal chain running from d_i to e will also constitute a *failure* in the causal chain running from c to e and will, therefore, be a stable weak decreaser for c and e such that there is either no increaser for c and d_j or no increaser for d_j and e.

Given these facts, Kvart's infinite regress condition can be replaced by the following, simple, universally quantified requirement: that for every stable weak decreaser, d, for c and e, there exists an increaser for c and d and an increaser for d and e. Making this replacement brings Kvart's condition into closer alignment with condition (ii) of my analysis.

6. Conclusion

In this paper I advanced a sophisticated probabilistic analysis of causation. This analysis gives the correct diagnosis of cases of non-probability-raising causation and cases of probability-raising non-causation. Unlike most existing probabilistic analyses, it does so in a manner that does not render it inconsistent with the possibilities of direct probability-lowering causation, causation by absence and omission, prevention, action-at-a-distance or intransitive causation. It therefore represents an improvement over these probabilistic analyses.

I observed that there is one sophisticated recent probabilistic analysis that comes closer to adequacy than the others. I showed that, though there are important similarities between this analysis and my own, where the analyses diverge there are good reasons for preferring the one advanced here.

Though I suggested at the end of §4 that my analysis deals better with certain difficult cases than do some of its major non-probabilistic rivals, I have not attempted a systematic evaluation of its merits vis-à-vis these (and other) rivals. I am optimistic that it would

compare favourably, but for now will be content if I have at least shown that the probabilistic approach is a viable option, and one that is worthy of further investigation.

References

Collins, John, Ned Hall and L.A. Paul (eds.) (2004): *Causation and Counterfactuals* (Cambridge, Mass.: The MIT Press).

Dowe, Phil (2000): *Physical Causation* (Cambridge: CUP).

Edgington, Dorothy (1997): 'Mellor on Chance and Causation', *The British Journal for the Philosophy of Science*, Vol. 48, No. 3 (Sept.), pp.411-433.

Eells, Ellery (1991): *Probabilistic Causality* (Cambridge: Cambridge University Press).

Good, I. J. (1961a): 'A Causal Calculus (I)', *The British Journal for the Philosophy of Science*, Vol. 11, No. 44. (Feb), pp.305-318.

Good, I. J. (1961b): 'A Causal Calculus (II)', *The British Journal for the Philosophy of Science*, Vol. 12, No. 45. (May), pp.43-51.

Good, I.J. (1962): 'Errata and Corrigenda', *The British Journal for the Philosophy of Science*, Vol. 13, No. 49 (May), p.88.

Hall, Ned (2004a): 'Causation and the Price of Transitivity' in Collins, Hall and Paul (eds.) (2004), pp.181-203.

Hall, Ned (2004b): 'Two Concepts of Causation', in Collins, Hall and Paul (eds.) (2004), pp.225-276.

Hesslow, Germund (1976): 'Two Notes on the Probabilistic Approach to Causality,' *Philosophy of Science*, Vol. 43, No. 2 (June), pp.290-292.

Hitchcock, Christopher (2001a): 'The Intransitivity of Causation Revealed in Equations and Graphs', *The Journal of Philosophy*, Vol. 98, No. 6 (June), pp.273-299.

Hitchcock, Christopher (2001b): "A Tale of Two Effects" *Philosophical Review* 110: pp.361-96.

Hitchcock, Christopher (2004): 'Do All and Only Causes Raise the Probabilities of Effects?' in Collins, Hall and Paul (eds.) (2004), pp.403-417.

Kvart, Igal (2004): 'Causation: Probabilistic and Counterfactual Analyses' in Collins, Hall and Paul (eds.) (2004). pp.359-386.

Lewis, David (1986a): *Philosophical Papers*, Vol. II (Oxford: OUP).

Lewis, David (1986b): 'Causation' in Lewis (1986a), pp.159-172.

Lewis, David (1986c): 'Postscripts to "Causation"' in his (1986a), pp.172-213.

McDermott, Michael (1995): 'Redundant Causation', *The British Journal for the Philosophy of Science*, Vol. 46, No.4 (Dec.), pp.523-544.

Menzies, Peter (1989): 'Probabilistic Causation and Causal Processes: A Critique of Lewis', *Philosophy of Science*, Vol. 56, No. 4. (Dec.), pp. 642-663.

Menzies, Peter (1996): 'Probabilistic Causation and the Pre-Emption Problem', *Mind*, New Series, Vol. 105, No. 417. (Jan.), pp. 85-117.

Pearl, Judea (2000): *Causality: Models, Reasoning and Inference* (Cambridge: CUP).

Reichenbach, Hans (1971): *The Direction of Time* (Berkeley and Los Angeles, California: University of California Press). (first published by University of California Press in 1956).

Rosen, Deborah A. (1978): 'In Defence of a Probabilistic Theory of Causality', *Philosophy of Science*, Vol. 45, No.4. (Dec.), pp.604-613.

Salmon, Wesley C. (1980): 'Probabilistic Causality', *Pacific Philosophical Quarterly*, Vol. 61, No. 1 (Jan.), pp.50-74.

Salmon, Wesley C. (1984): *Scientific Explanation and the Causal Structure of the World* (Princeton, N.J.: Princeton University Press).

Suppes, Patrick (1970): *A Probabilistic Theory of Causality*, *Acta Philosophica Fennica* XXIV (Amsterdam: North-Holland Publishing Company).

Woodward, James (2003): *Making Things Happen* (Oxford: OUP).