First publ. in: Logic, Language, and the Structure of Scientific Theories : proceedings of the Carnap-Reichenbach Centennial, University of Konstanz, 21 - 24 may 1991/ ed. by Wesley Salmon ... Konstanz : Univ.-Verl. Konstanz, 1994, pp. 215-239

On Reichenbach's Principle of the Common Cause[*]

Wolfgang Spohn

Abteilung Philosophie
Universität Bielefeld
4800 Bielefeld 1
Germany

This paper deals with Hans Reichenbach's common cause principle. It was propounded by him in (1956, ch. 19), and has been developed and widely applied by Wesley Salmon, e.g. in (1978) and (1984, ch. 8). Thus, it has become one of the focal points of the continuing discussion of causation.

The paper addresses five questions. Section 1 asks: What does the principle say? And section 2 asks: What is its philosophical significance? The most important question, of course, is this: Is the principle true? To answer that question, however, one must first consider how one might one argue about it at all. One can do so by way of examples, the subject of section 3, or more theoretically, which is the goal of section 4. Based on an explication of probabilistic causation proposed by me in (1980), (1983), and (1990), section 4 shows that a variant of the principle is provable within a classical framework. The question naturally arises whether the proved variant is adequate, or too weak. This is pursued in section 5.

My main conclusion will be that some version of Reichenbach's principle is provably true, and others may be. This may seem overly ambitious, but it is not. The paper does not make any progress on essential worries about the common cause principle arising in the quantum domain; it only establishes more rigorously what has been thought to be plausible at least within a classical framework.

1. What does the principle say?

The principle of the common cause specifies an important relation between probability and causality. Though requiring some explanation, its statement is straightforward: Let A and B be two positively correlated events, i.e. events which satisfy the condition

(1) $P(A \cap B) > P(A)\,P(B)$.

Then one of the events causes the other or there is a further event C which is a *common cause* of A and of B such that

(2) $P(A \mid C) > P(A)$, and $P(B \mid C) > P(B)$, i.e. C is positively relevant to A and to B, and

(3) $P(A \cap B \mid C) = P(A \mid C)\,P(B \mid C)$, and $P(A \cap B \mid \overline{C}) = P(A \mid \overline{C})\,P(B \mid \overline{C})$, i.e. A and B are independent conditional on C and on $\overline{C}$.

(2) and (3) indeed imply (1); and given the independencies in (3), (2), with the inequalities reversed, would also imply (1), whereas (2), with only one inequality reversed, would imply the reversal of (1).

This formulation slightly generalizes Reichenbach's original statement in a way suggested by Salmon (1980, 61). The original principle is obtained by assuming additionally that A and B occur simultaneously in which case the principle leaves no alternative to a common cause of A and B because there cannot be a causal influence running from A to B or the other way around. But what does all of this really mean?

Since the common cause principle seems to be about events (though I shall argue in section 5 that it should rather be viewed as being about (random) variables), something must to be said about what an event is. Obviously it must be something which serves as an object of probability *and* as a causal relatum. By definition, objects of probability are events in the mathematical sense; that's just mathematical usage, and perhaps inappropriate from a philosophical point of view. There is, however, no common positive philosophical view. The predominant one seems to be that an event is a concrete particular located in space and time, which is in fact individuated by its spatio-temporal boundaries, and which could not have been realized in any other way.[1] However, such concrete particulars cannot be objects of probability, simply because logical operations like negation, disjunction, etc., can be applied to objects of probability, but not to concrete particulars.[2] This entails in particular that events in the latter sense cannot be relata of probabilistic causation.

This negative conclusion is important, I think, and raises an issue about how really to understand events. Since it cannot be fully addressed here, let me say only that throughout the paper events are taken to be a kind of propositions or states of affairs, namely singular, temporally located propositions which typically consist in a certain object having a certain property at a certain time. Events in this sense seem suited as objects of probability and as causal relata. This, in fact, is how Kim (1973), for instance, explicates events.[3]

However, this is not quite how Reichenbach understood events. The deviation is related to the next point of clarification: What is probability? This is an even bigger issue which I would like to leave open here. An objective interpretation of probability may seem preferable in the given context. But a subjective interpretation seems to me to be entertainable as well, although the common cause principle should then be taken as speaking about the causal conception of the subject at hand.[4] If probabilities *are* taken objectively, there is again an alternative. It is preferable to view them as some kind of propensities because these apply to singular events. Reichenbach, however, adopted a frequentistic conception of objective probability, and a peculiar one at that. Hence, his formulation of the principle referred rather to generic events or event types. I want to leave aside this specific part of Reichenbach's doctrine. Therefore I will stick to singular events instead of event types as objects of probability; this is to be the only constraint on the interpretation of probability.

Is the principle to be viewed as implicitly containing a definition of common causes? That is, is *C* a common cause of *A* and *B* if and only if (2) and (3) hold true? Reichenbach is somewhat vague on that point[5], and by introducing so-called interactive forks Salmon

(1978) argues that there may be other kinds of common causes. In any case, it is more cautious to take (2) and (3) not as a definition, but just as a necessary condition on the common cause.

But even on that weaker interpretation the demand arises: Why (2) and (3) and not any other conditions? Concerning the positive relevance conditions in (2), the answer is that the positive probabilistic relevance of a cause to its effect has almost without exception been taken as the minimal core of any probabilistic treatment of causation.[6] Indeed, this justifies also the first alternative of the common cause principle: the positive relevance expressed by (1) may also be directly due to a causal relationship between *A* and *B*.

Concerning the conditional independence conditions in (3), footnote 6 quotes a reason actually given by Reichenbach; but it is more fully stated by van Fraassen. It is simply that when the correlation between *A* and *B* is also conditional on *C* or on $\overline{C}$, there remains the need of explaining that residual correlation and thus of finding further common causes, and this continues to be so until an event (or rather a partition of events) is found which renders *A* and *B* conditionally independent.[7] Hence, what Reichenbach's principle in effect postulates is the existence of a *total* common cause (which, by the way, is very different from a deterministic or even sufficient common cause). Common causes in general, it seems, need not satisfy the condition (3); but if they are to be total, (3) is obligatory.

This may suffice as an explanation of what the principle says. But it also has philosophical implications which should be made a bit clearer before proceeding any further.

2. What is the philosophical significance of the principle?

The most important thing to note is that the common cause principle is a descendent of the principle of causality which says that each event has a cause. Both principles claim the existence of a cause. One might say that the one claims the existence of causes for events and the other does so for correlations. This is misleading, however, because correlations may be said to have causes at most in a derivative sense. The difference is, more accurately, that the principle of causality postulates causes in full generality, whereas the common cause principle, though postulating *common* causes, does so only in the restricted case of a positive correlation between two causally unrelated events.

In view of this kinship it is not surprising that both principles are subject to nearly the same variation in attitudes. Both principles have been taken not as making any claims at all, but as having a practical function in guiding research and in encouraging scientists not to give up hope in their search for causes. Both principles have been taken as making empirical assertions which are presumably wrong. However, both principles have also been thought to have some kind of a necessary status. Perhaps they turn out to be analytically true on the basis of an adequate definition of events and/or of causation; indeed, this is my attitude toward the common cause principle, as I will explain below.[8] Or perhaps they have another kind of necessity; e.g., they may be knowable a priori, a view I think in a way to be the case with the principle of causality.[9]

Besides the kinship there is a further remarkable difference between the principles. The old principle of causality has always been conceived in a deterministic way; causes tended to be taken as sufficient causes. By way of contrast, the common cause principle emphasizes a probabilistic conception of causation. Doubtlessly, this latter was a pioneering insight of Reichenbach.

How important this emphasis is may best be seen by looking at statistics. Statisticians typically are ambivalent. On the one hand, the desire to discover causal relationships has always been an essential motive in doing statistics; but on the other hand, statisticians have always been unsure about how to infer causal from statistical relations. The common cause principle in itself does not tell which statistical correlations to interpret as causal ones. Nevertheless, it certainly serves a crucial bridging function by specifying a connection between probability and causality.

However, this was not the immediate interest of Hans Reichenbach. His concern was about a different and deeply philosophical issue, namely the causal theory of time. His idea was the following: If *A*, *B*, and *C* satisfy (1) - (3), let's say that they form a conjunctive fork which is closed at *C* and which may be open or closed at the other side according to whether there is an event *D* on the other side standing to *A* and *B* in the same probabilistic relation as *C*. Reichenbach's ingenious contention was that any two conjunctive forks which are not closed on both sides are open on the same side. Then we may define the direction of time in a purely probabilistic way by stipulating that the open forks are open towards the future. Of course, it's easy to find purely numerical counter-examples. But as far as I know, no convincing physical counterexample has turned up in the literature. So Reichenbach's contention still stands.

I shall add only one skeptical remark here. The idea looks circular because objective probability makes reference not only to time, but also to the direction of time. On a propensity view, the chance of an event may change over time until it reaches the event's "end-point chance", as Lewis (1980, 271) puts it; afterwards it's 1 or 0. That is how chance is relative to time. Moreover, the chance of an event at a certain time depends only on all of its past, but not at all on its future, as has been made clear by Lewis (1980, 272f.). That's how chance refers to the direction of time, and, thus, the reduction of time to chance seems spurious.

This argument starts from a propensity interpretation of probability. But it may have an echo in the frequentist account. On such an account we have to count relative frequencies, and we must start counting somewhere in the middle of the world's total time span. Whether or not one or both directions of time are infinite, the relative frequency or its limit obviously depends on the direction in which we count. Thus it seems that even such a concept as limiting relative frequency presupposes the direction of time.

A further philosophical use of the common cause principle should finally be mentioned. Several philosophers have argued that another important principle, the inference to the best explanation, supports a realistic attitude as opposed to attitudes of an idealist or empiricist brand. The argument is simply that realism *is* the best explanation of our sensations, beliefs and similar things of which we are immediately aware.[10] If this is true, then the common cause principle seems to provide even stronger support for scientific realism, because it starts not only from one, but from two givens: it does not infer an explanans from *one* explanandum as does the inference to the best explanation; rather it proceeds from *two* supposed effects and infers the existence of a usually hidden common cause, one which need not be located at the same level of observationality as the effects. To paraphrase a trigonometrical metaphor of Donald Davidson, with a twist towards philosophy of language: The members of a speech community triangulate reality from their very parallel behavior which would otherwise be a most surprising coincidence; and the triangulation proceeds precisely via the common cause principle.

In fact, this issue is the deeper motive of the debate between Wesley Salmon and Bas van Fraassen; the one defends a weakened version of the common cause principle in support of his realistic attitude (most expressly in Salmon (1984, ch. 8)), whereas the other criticizes it in order to promote his empiricistic project (started in van Fraassen 1980). I will not attempt to comment on that debate or on the argument from the common cause principle to realism. But it is important to keep in mind the philosophical fundamentals which are here at stake.

3. How might one argue about the principle?

Having exposed the content and the significance of the common cause principle, I now want to take it at face value, namely as making a claim. Usually the most interesting question concerning a claim is whether it is true. In the present case it's hard to say yes or no; and it's even hard to know how to go about establishing or refuting the claim. So the initial question concerns how one might argue about the principle at all.

There are essentially two ways: One may either study particular examples and applications and see whether they confirm or disconfirm the principle. Or one may go more deeply into the theory of causation and see whether illumination of a more theoretical kind is forthcoming. I prefer the second way and shall ultimately stake my case there. But let's briefly look at the first alternative.

There are many instantiations of the common cause principle. But because of mixed quantifiers it is much more difficult to find good counter-examples. What one would have to show somehow is that there is no common cause in a given case. Surprisingly, at least one striking counter-example has been found. I refer to the situations as they are conceived in the Einstein-Podolski-Rosen paradox and as developed by D. Bohm and J.S. Bell.

In these situations, a pair of elementary particles, say an electron and a positron, are prepared in the so-called singlet spin state. From their common source, one is emitted to the right and the other to the left. After some time, they simultaneously meet two Stern-Gerlach magnets which deflect them downwards or upwards and thus measure their spin in the direction of the magnets. What exactly happens either on the left or on the right is up to chance; in either case, the chance of the particle being deflected downwards or upwards is, respectively, 1/2. But nevertheless, a correlation between the left and the right particle being deflected into different directions may be observed. If the magnets are parallel, this correlation is perfectly positive, i.e. 1; whenever one particle goes down, the other goes up. If the magnets are not parallel, there is still a positive correlation depending on the angle at which the magnets have been rotated; and that dependence is precisely described by quantum mechanics.

So, the common cause principle should apply to such situations. And since the right and the left event are simultaneous, one should find a common cause for them satisfying the conditions (2) and (3) above. Now, whenever the joint distribution for the left and the right side is generated by a distribution for a third variable conditional on which the left and the right side are independent, then that joint distribution must satisfy a certain inequality named after J.S. Bell. However, and this is the upshot, this inequality is demonstrably violated by the joint distribution derivable from quantum theory. Hence, if quantum theory is right, there is demonstrably no hidden variable, no common cause accounting for the correlations obtaining in such situations. This is, very briefly, the most powerful argument against Reichenbach's common cause principle.[11]

The quantum phenomena, thus, seem to provide the central test in which the common cause principle fails. On the basis of this set-up, van Fraassen (1982b) and Salmon (1978) and (1984, ch. 6) have devised a number of further counter-examples, with diverging intentions. Salmon wants to save the common cause principle by weakening it to the effect that the common cause and its two effects may form either a conjunctive fork, as required by Reichenbach, or a so-called interactive fork, where "=" is replaced by ">" in the first equation of (3); this liberalization is, however, not intended to cope with the original EPR phenomena. Van Fraassen, on the contrary, wants to generally discredit the common cause principle, degrading it to "a tactical maxim of scientific inquiry and theory construction. ... its acceptance does not make one irrational; but its rejection is rationally warranted as well" (1982b, 209).

However this may be, in the sequel I shall consider the common cause principle only within a classical framework. For, in this domain, there seem to be no convincing counter-examples to the principle. But is lack of convincing counter-examples sufficient evidence for accepting the principle on classical terms? What is really needed is a kind of explanation why the principle is or should be true.

4. How else might one argue about the principle?

So let us try to get a more theoretical perspective on probabilistic causation and hence on the common cause principle. More specifically, I want to determine the status of the principle within the theory of causation I proposed in (1980), (1983), and (1990) and to work up to the variant of it which is provable on the basis of that theory. This theory keeps within the mainstream of the relevant literature, but differs from other accounts in some respects.

Each discussion of probabilistic causation must proceed from an explicitly given probability space: So let I be a non-empty set of variables or factors; I call it a *frame*. Each variable $i \in I$ is associated with a set Ω_i of at least two possible values i may take. Ω is to be the cross product of all the Ω_i, i.e. the set of all functions ω defined on I such that for each $i \in I$, $\omega(i) \in \Omega_i$; intuitively, each ω represents a possible course of events - a possible world in philosophers' talk, or a possible path in the mathematicians' terminology.

To give just one example: In meteorology, one is interested in several items: temperature, atmospheric pressure, humidity, precipitation, cloudiness, velocity, direction of the wind, etc. For each place and time considered these items constitute separate variables. A possible world, as far as the meteorologist is concerned, consists then in a specification of the values of all these variables.

Essentially for the sake of mathematical simplicity I assume that I, each Ω_i, and hence Ω are finite. In particular this implies that questions of measurability can be neglected because each subset of Ω can be taken to represent an event in the mathematicians' sense, or a proposition in the philosophers' sense. Further, I assume a probability measure P assigning a probability to each proposition, i.e. to each subset of Ω. This completes the description of the underlying probability space.

I shall assume that the probability measure P is strictly positive, i.e. that $P(\{\omega\}) > 0$ for all $\omega \in \Omega$; hence, the conditional probability $P(B \mid A)$ is defined for each $A \neq \emptyset$. Since Ω is finite, this assumption is unproblematic. The reason is that all probabilistic theories of causation run into serious problems with the limiting probabilities 0 and 1; and since these problems are not germane to the present topic, they can be excluded without prejudice.

Next, the possible worlds must be provided with a temporal structure. For this purpose, I assume a weak, i.e. transitive and connected order relation $\leq$ on the set I of variables which represents the order of the times at which the variables are realized; $<$ is to denote the corresponding irreflexive order relation. Some abbreviations will be useful: For $j \in I$ and $K \subseteq I$ we put $\{< j\} = \{k \in I \mid k < j\}$ and $\{< j - K\} = \{< j\} - K$; $\{\leq j\}$ and $\{\leq j - K\}$ are defined correspondingly.

The temporal order may be partially extended to propositions in the following way. For any $J \subseteq I$ define a proposition A to be J-measurable or a J-proposition iff for all ω, $\omega' \in \Omega$ agreeing on J, $\omega \in A$ if and only if $\omega' \in A$; intuitively, a J-proposition refers at most to the variables in J and not to any variable outside J. In particular, there are propositions about single variables, and it is them to which the temporal order can be immediately carried over. I shall take only such propositions referring to single variables as *events*, i.e. as causal relata; there is no need to consider logically complex propositions as causes or effects.

For $\omega \in \Omega$ and $J \subseteq I$ we shall have to consider the proposition $\{\omega' \mid \omega'(i) = \omega(i)$ for all $i \in J\}$ which says that the variables in J behave as they do in the world ω; it will be denoted by $^{\omega}J$. Obviously, this proposition is a J-proposition. In fact, it is an atom of the algebra of J-propositions. I call it a *J-state*.

Finally, we need a notation for (conditional) probabilistic independence. $A \perp B$ says that the propositions A and B are probabilistically independent, i.e. that $P(A \cap B) = P(A)\, P(B)$; and $A \perp B / C$ says that A and B are independent conditional on C, i.e. that $P(A \cap B \mid C) = P(A \mid C)\, P(B \mid C)$. These notions can be extended to variables and sets of them. For $K, L, M \subseteq I$ $K \perp L / M$ means that $A \perp B / C$ for all K-propositions A, L-propositions B, and M-states C; and it will be clear what mixed formulae like $K \perp A$ or $K \perp L / C$ are to mean.

The basic notion may now be introduced in full rigour. Informally, we may say that A is a cause of B if and only if A and B both obtain, if A precedes B, and if A raises the epistemic or metaphysical rank of B under the obtaining circumstances. This is the basic conception of causation on which most can agree. In the deterministic case, it covers regularity theories, counterfactual approaches, and analyses in terms of necessary and/or sufficient conditions; the difference is only on the relevant meaning of "raises the epistemic or metaphysical rank". In the probabilistic case, this can only mean one thing, namely that A raises the probability of B, where, however, probabilities can be interpreted subjectively or objectively.

The condition that causes and effects have to obtain makes the notion of causation world-relative; in our framework we may obviously say that A and B obtain in ω iff $\omega \in A \cap B$. The condition that A precedes B is expressed by the clause that there are variables i and j such that A is an i-proposition, B a j-proposition, and $i < j$; but a bit more needs to be said about this later on. The condition that A raises the probability of B sounds as if A would do something; but it just means that the probability of B given A is larger than given $\bar{A}$. Finally, the phrase "under the obtaining circumstances" is beset with great difficulties. But if we restrict ourselves to direct unmediated causation, there is a particularly simple explanation. As I have argued in the papers referred to, each fact preceding the direct effect B and differing from the direct cause A is to count among the obtaining circumstances of the direct causal relation between A and B; whenever judg-ment about that relation is based on less, it may be just the neglected facts which would change the judgment. This means that in the world ω the obtaining circumstances consist of the whole past of B in ω with the exception of A; in the above notation this is just the proposition $^{\omega}\{< j - i\}$. So we arrive at the following explication:

A is a direct cause of B in ω iff $\omega \in A \cap B$ and there are variables $i, j \in I$ such that A is an i -proposition, B a j -proposition, $i < j$, and $P(B \mid A \cap {}^{\omega}\{< j - i\}) > P(B \mid \bar{A} \cap {}^{\omega}\{< j - i\})$. It is natural to add the definition that *A is a direct counter-cause of B in* ω iff $\omega \in A \cap B$ and there are variables $i, j \in I$ such that A is an i-proposition, B a j-proposition, $i < j$, and $P(B \mid A \cap {}^{\omega}\{< j - i\}) < P(B \mid \bar{A} \cap {}^{\omega}\{< j - i\})$. Thus, *A is directly causally relevant to B in* ω iff A is a direct cause or counter-cause of B in ω, i.e. iff not $A \perp B \,/\, {}^{\omega}\{< j - i\}$.

How is this to be extended to a general account of causation? This is a complicated story. The view which I arrive at in (1990), and which I simply assume here, is that *causation in general* should be defined as the transitive closure of direct causation. This had already been suggested by Lewis (1973) in his counterfactual explication of deterministic causation.

It may be illuminating to pause here and make some brief remarks comparing this explication with accounts of probabilistic causation given by Suppes (1970), Good (1961-63), and Cartwright (1979).

Starting from his definition of prima facie causes which does not refer to any given background Suppes (1970) considers the circumstances of causal relationships for distinguishing spurious versus genuine and direct versus indirect causes. On pp.41f. he acknowledges the legitimacy and usefulness of relativizing all his definitions to some background information. However, one has moreover to distinguish between overt and hidden causes; and as I have argued in (1980) and (1983), the interplay of these three distinctions forces one to consider much richer backgrounds, and indeed different backgrounds for different causal relationships. Thus I arrived at the above definition which explicitly refers to the richest possible circumstances of direct causal relationships.

The theory of Good (1961-63) differs from my explication in several ways, but the crucial point is that in defining the tendency of *A* to cause *B* Good considers different conditional probabilities. In (1961, 308f.) he conditionalizes on the whole past of the cause and on all laws of nature, whereas I conditionalize on the whole past of the direct effect. I have not found a clear argument for the appeal to the laws. In fact, I think it diminishes the philosophical interest of the project, since one may hope that an analysis of singular causation will further our understanding of laws of nature.[12] The main question, however, is whether to conditionalize on the past of the cause or on the past of the effect. If one wants to have one explication of both, direct and indirect causation, then Good's conditionalization policy seems certainly more plausible. But, as I argued in (1990, 128ff.), Good's account is an inadequate account of direct causes. This was one of the reasons why I split up the explication into a more adequate definition of direct causation and its extension to indirect causation.

Cartwright (1979) is interested rather in causal laws than in singular causation. Still, it is instructive to compare her views with my explication. In a way, she explains forcefully why Simpson's paradox is a crucial problem for probabilistic theories of causation. And in a way, my explication proposes a radical solution; if one conditionalizes on the whole past of the effect, then there is no further subdivision of that past which could change the conditional probabilities, at least within the descriptive frame given by the set *I* of variables. But this is not her solution. She argues that all the variables influencing *B* but not influenced by *A* constitute the obtaining circumstances of the causal relation between *A* and *B* and that it is conditionalization with respect to these variables which indicates whether *A* is a cause of *B*.

The disagreement is less substantial than it seems, however. Cartwright rightly insists that one must not conditionalize with respect to variables mediating between cause and effect; indeed, if their values are given, the cause can no longer be expected to be positively relevant to the effect. But in the special case of direct causation there are no mediating variables; and the difference then reduces to the fact that I conditionalize also with respect to all variables which precede, but do not influence the effect, whereas she does not. I believe the more extensive conditionalization proposal to be harmless, but she does not. She says that "partitioning on an irrelevancy can make a genuine cause look irrelevant, or make an irrelevant factor look like a cause" (1979, 432) and goes on to illustrate this alleged possibility. But I do not think that this illustration supports her restricted form of conditionalization, as Eells and Sober (1983, 42), who also conditionalize on irrelevant factors, have already argued.

In one respect, however, the disagreement is deeper. Cartwright's conclusion is that there is no non-circular characterization of causation in probabilistic terms. My conclusion, on the contrary, is again that direct and indirect causation should be considered separately. My explication of direct causation is in line with her ideas and not threatened by circularity. From this fact one can proceed to deal with the circularity involved in the explication of indirect causation; and if my 1990) account of causation in general is tenable, this circularity can also be dissolved.

The preceding may suffice as a setting of the above explication within a family of related views. A final remark is necessary concerning the requirement that the cause temporally precedes the effect.

Though some philosophers and physicists would like to make sense of the possibility that the effect is later than the cause, I am utterly skeptical. In any case it should be clear that this possibility is very foreign to the present approach. But there remains the question whether to take precedence strictly or loosely so as to include simultaneity. Here we have choice. One alternative is to adapt my approach to the possibility of simultaneous causation. This is easily done; simply replace strict precedence by precedence or simultaneity throughout the above explication. Such was my choice in (1980). The mathematics then yields many desired theorems and also the one to be proved below. The disadvantage, of course, is that the only causal relation between simultaneous events is interaction; for them there is no way to tell cause and effect apart.

The other alternative is to stick to a strict interpretation of precedence. Then, however, we run into mathematical problems. The desired theorems no longer follow in full generality; special assumptions are required. Only two such assumptions are sufficient, as far as I can see. One is to assume that there are no simultaneous variables so that the temporal order is turned from a weak into a linear order. When spatial considerations are added within a relativistic framework, this alternative may be attractive because the signal relation yields a strict order relation between space-time points and thus between variables realized at these points. But within the present framework, which is intended to be much more widely applicable, this assumption is a severe limitation. The alternative assumption is a kind of locality condition, namely that all simultaneous variables are independent given all of their past. If $\{\approx i - i\}$ denotes the set of all variables in I simultaneous with i except i itself, the condition is, formally:

(L) $\quad i \perp \{\approx i - i\} \,/\, (< i\}$ for all $i \in I$.

This assumption (L) may be interpreted as a condition on probabilities or as governing our understanding of what counts as a single variable. It seems to be characteristic of a classical framework that in each case a frame I violates (L) the suspicion arises that that frame is incomplete and should be completable so as to satisfy (L). Indeed, a salient feature of the quantum theoretical description of the EPR situations is that they violate (L).

Note that if there are no simultaneous variables, (L) is vacuously true. Thus, in either case (L) will suffice to prove my version of the common cause principle. This is not question-begging. True, (L) says that simultaneous events are independent under certain conditions; but (L) by itself does *not* say that there exist common causes or that they form such conditions.

What is the version of the common cause principle to be proved? To state it, the causal notions so far introduced have to be extended to variables. This is straightforward. For two variables i and j, *i is directly causally relevant to j in* ω iff some i-proposition is directly causally relevant in ω to some j-proposition, that is, iff not $i \perp j \,/\, ^{\omega}\{< j - i\}$; *$i$ is potentially directly causally relevant to j* iff i is directly causally relevant to j in some world, i.e. iff not $i \perp j \,/\, \{< j - i\}$; general, i.e. direct *or* indirect *causal relevance in* ω is again just the transitive closure of direct causal relevance in ω; and likewise for *potential causal relevance*. Here, "potential" refers roughly to the causal relations within possible courses of events which need not be the actual one. This is not merely a logical or some other vacuous possibility; it is substantially restricted by the given frame I and the given probability measure P. This point will be of some importance later on.

Now the theorem to be proved says: Let i and j be any two variables in I, $K = \{k \in I \mid k = i$ or k is potentially causally relevant to $i\}$, and $L = \{l \in I \mid l = j$ or l is potentially causally relevant to $j\}$. Then, given condition (L), $K - L \perp L - K \,/\, K \cap L$ and, a fortiori, $i \perp j \,/\, K \cap L$ hold true.

The proof of this theorem relies crucially on the laws of conditional independence; indeed, they lie at the very heart of any probabilistic theory of causation. I state the most important ones (there are more) without proof.[13] For all $J, K, L, M \subseteq I$ we have:

(I) if $K \perp L \,/\, M$, then $L \perp K \,/\, M$;

(II) if $K \subseteq M$, then $K \perp L \,/\, M$;

(III) if $K' \subseteq K \cup M$, $L' \subseteq L \cup M$, $M \subseteq M' \subseteq K \cup L \cup M$, and $K \perp L \,/\, M$, then $K' \perp L' \,/\, M'$;

(IV) if $J \perp K \,/\, L \cup M$ and $J \perp L \,/\, M$, then $J \perp K \cup L \,/\, M$;

(V) if K and L are disjoint, $J \perp K \,/\, L \cup M$, and $J \perp L \,/\, K \cup M$, then $J \perp K \cup L \,/\, M$, provided P is strictly positive (as was assumed here).

Proof of Theorem:

Let i, j, K, and L be as in the theorem; we may assume that $i \leq j$. Let $\leq^*$ be any linear order on I agreeing with the temporal order $\leq$; i.e. $k < l$ entails $k <^* l$. The notation $\{\leq^* k\}$, $\{<^* k\}$, etc. is explained as before. We now show inductively that for each $m \leq^* j$

(4) $K \cap \bar{L} \cap \{\leq^* m\} \perp \bar{K} \cap L \cap \{\leq^* m\} \,/\, K \cap L \cap \{\leq^* m\}$.

If m is the first variable relative to $\leq^*$, (4) is trivially true. For the induction step we assume that

(5) $K \cap \bar{L} \cap \{<^* m\} \perp \bar{K} \cap L \cap \{<^* m\} \,/\, K \cap L \cap \{<^* m\}$

and shall show that (4) follows. For this purpose four cases must be distinguished:

First, if $m \notin K \cup L$, then (4) immediately reduces to (5) because all the terms in (4) are identical with those in (5).

Secondly, assume that $m \in K \cap \bar{L}$. Then we have for all $k \in \bar{K} \cap \{< m\}$

(6) $m \perp k \,/\, \{< m - k\}$,

because each such k is not potentially directly causally relevant to m. (6) entails, with the help of (V),

(7) $m \perp \bar{K} \cap \{< m\} \,/\, K \cap \{< m\}$.

(7) and (L) entail, with the help of (IV) and (III),

(8) $m \perp \bar{K} \cap \{<^* m\} \,/\, K \cap \{<^* m\}$.

(III) allows to weaken (8) to

(9) $m \perp \bar{K} \cap L \cap \{<^* m\} \,/\, K \cap \{<^* m\}$.

According to (IV), (5) and (9) entail

(10) $K \cap \bar{L} \cap \{\leq^* m\} \perp \bar{K} \cap L \cap \{<^* m\} \,/\, K \cap L \cap \{<^* m\}$.

And (10) is equivalent to (4), since their second and third terms are identical.

The third case that $m \in \bar{K} \cap L$ is symmetrically similar to the second case.

In the final case when $m \in K \cap L$ (6) - (10) hold as well; with (II) (10) implies

(11) $K \cap \bar{L} \cap \{<^* m\} \perp \bar{K} \cap L \cap \{<^* m\} \,/\, K \cap L \cap \{\leq^* m\}$;

and now (11) is equivalent to (4).

Hence (4) holds for all $m \leq^* j$; and for $m = j$ (4) is the assertion to be proved.

As to the final assertion that $i \perp j \,/\, K \cap L$: If $i \in K \cap \bar{L}$ and $j \in \bar{K} \cap L$, it follows immediately from what has just been proved; and if i or j is in $K \cap L$, it is trivially true. This completes the proof.[14]

This theorem is surely a variant of the common cause principle because it entails two facts: First, if the i-proposition A is positively or negatively correlated with the j-proposition B, then $i \perp j$ does not hold. This entails according to the theorem that $K \cap L$ is not empty; that is, either the earlier of i and j is in $K \cap L$ and hence potentially causally relevant to the later one, or there is at least one variable potentially causally relevant to both. This is one part of the common cause principle.

Secondly, though the theorem does not confirm that i and j are independent given some common cause, it says that they are independent given the set of all the variables potentially causally relevant to both, i and j, which may well be said to form a total common cause in each world; and this is, in a way, the other part of the common cause principle.

This apparently is a close approximation to Reichenbach's common cause principle given the above explication of probabilistic causation. If so, then if one is convinced of the common cause principle anyway, the basis on which it is reconstructed is thereby rendered more plausible. And if the basis is plausible, the derivation presents a good explanation of the validity of the common cause principle. This is surely true in a classical framework as characterized at least by the condition (*L*); concerning the quantum theoretical context our understanding of causation and of the role of the common cause principle is certainly not enhanced.

However, it is too early to summarize because the theorem also diverges from Reichenbach's original principle. The question arises, therefore, to what extent is the presented approximation really satisfactory?

5. Is the proved variant too weak?

There are four points of divergence from the original. First, condition (2) does not seem to find a counterpart in the proved variant. Secondly, whereas the original claims the existence of some common cause, the variant refers to the total common causal ancestry. Thirdly, the variant assumes two correlated variables instead of two positively correlated events, and the common cause is also replaced by a set of variables. Finally, wherever the original talks of causation, the variant talks of potential causal relevance. I shall take up these points in that order.

First, concerning condition (2), their counterparts in the terminology of the above theorem are given by the assertion that neither i nor j is independent of $K \cap L$. This should be provable, it might seem, but it is not; lucky averaging may create surprising independencies. The theorems 14 and 16 of my (1990) state assumptions under which also an indirect cause is conditionally positively relevant to its effects (this was assumed in the above explication only for direct causes), or rather assumptions under which probabilistic dependence spreads through chains of causal relevance. With their help one may specify when i and j are also probabilistically dependent on their common causal ancestry.

Secondly, the fact that the theorem requires reference to the total common causal ancestry of the correlated variables confirms the conjecture at the end of section 1 that in his principle Reichenbach thinks of a common cause as a total one; otherwise the conditional independence of the correlated variables might well fail.[15] Still, the reference to the total common ancestry, however remote, is not really necessary; what is required is only a complete cross-section, so to speak, of that total ancestry. Thus, the above proof can easily be adapted for showing that in the theorem $K \cap L$ can be replaced by the subset $\{k \in K \cap L \mid$ there is an $l \in (K \cup L) - (K \cap L)$ such that k is potentially directly causally relevant to $l\}$, i.e. the most proximate part of the total common causal ancestry. This is a slightly more satisfying result.

Thirdly, why should one refer to variables rather than to events? This is related with the switch from causation to causal relevance. Both changes seem to be required:

If there are positively correlated events, there also are negatively correlated events. Surely, we would like to have a causal explanation of that negative correlation as well and expect there to be one, if not necessarily in terms of common causes, at least in terms of common causally relevant factors.

Or consider again (1) - (3). I have mentioned that (2) and (3) imply (1) also when the inequalities in (2) are reversed. In that case *C* would be rather a common counter-cause than a common cause of *A* and *B*. And this may certainly happen within a probabilistic context where one must reckon for counter-causation no less than for causation (cf. footnote 7). Probable events may occur or not; and a cause makes its effect only more likely, not necessary, and its counter-effect only less likely, not impossible. Thus, the common cause principle should be taken not as excluding situations with the inequalities in (2) reversed, but as covering them as well.

Or consider a variation of the theme. Suppose that the events *A*, *C*, and *B*, thus temporally ordered, occur and form a Markov chain, i.e. *B* is probabilistically independent of *A* given *C* as well as given $\overline{C}$. Suppose further that *A* is in fact a counter-cause of *C* and *C* a counter-cause of *B*. Thus, the probabilistic side of the situation is again descri-bed by (3) and the reversed (2) Hence, (1) holds as well, i.e. *A* and *B* are positively correlated. The point is that this situation violates the original principle, since *A* clearly is not a cause of *B* (counter-cause plus counter-cause does not add to a cause) and since we may assume that there is no third event causing *A* and *B*. Again, the conclusion is, I think, that we should take the principle so liberally as to include that situation; and if we retreat to causal relevance, we do so, since *A*, though not causing *B*, is causally relevant to it.

The general lesson is this: Within the probabilistic realm, counter-causation or negative causal relevance is as ubiquitous as causation or positive causal relevance; things almost always happen because of some circumstances and despite of other circumstances. Moreover, in complex causal nets positive and negative causal influences may mix in countless ways. If we want to take account of all this, we have to attend to causal relevance *simpliciter*. Of course, it is important to disentangle the various kinds of causal relevance, but not in the present context where we have to consider all of them.

Now the natural relata of causal relevance are variables (and sets of them). The reason is this: If one event is positively relevant to another, that relation is turned negative by negating one of the events and positive again by negating both (this holds for causal as well as for probabilistic relevance). Thus, positive and negative relevance pertains to events and only to them. But if we are concerned with relevance only, it does not matter whether we take events or their negations. But then, in effect, we consider binary variables; and if we do so, the natural step is to generalize and to consider arbitrary variables.

This answers already part of the final point; the principle should indeed talk not of causation, i.e. positive causal relevance, but of causal relevance *simpliciter*. However, causal relevance may be actual or potential. Above, both kinds of relevance have been defined, but the theorem focused on the less plausible alternative, i.e. on potential relevance. There is a mathematical reason: it is open whether something similar can be proved on the basis of actual relevance. And there is a conceptual reason: if we really want to turn to variables, we also have to turn to potential relevance, because only that relation refers only to variables; actual causal relevance is conditioned on a proposition, i.e. the actual past. Indeed, in many contexts only variables and sets of them are considered when discussing causal matters, as is exemplified, e.g., in Granger (1980), Kiiveri et al. (1984), Glymour et al. (1987), and Pearl (1988).[16]

However, the switch to potential relevance may seem to reduce the impact of my variant of the common cause principle. In particular, it may seem that it is vacuously true because for any two variables *i* and *j*, whether they are correlated or not, there is a third variable *k* which is potentially causally relevant to both, i.e. which is causally relevant to *i* in some world and to *j* in some other world.[17] This is easily imaginable since one can invent arbitrary causal connections in possible worlds. Surely Reichenbach cannot have meant this triviality.

One must observe, however, what potential causal relevance means. Possibility is here constrained in two ways. First, it is constrained by the given frame *I*. A possible world is taken here, and elsewhere, just as a realization of the variables in *I*; I do not allude to the understanding of worlds as spatio-temporally (and causally) maximally inclusive things. Hence, possible causal connections cannot simply be invented by assuming new variables and suitable realizations of them.

Now one may wonder about the frame-relativity of the given explication of causation. The question whether this contains an important truth about causation or whether it renders the explication inadequate, is a deep one which I shall not fully discuss here.[18] But I will say this: The only way to eliminate that relativization in the present framework is to refer to a fictitious universal frame *I** just rich enough to completely describe our actual world. But even *I** is not large enough to generate all possible worlds whatsoever. The generated possible worlds are strictly made out of the material of the present world, so to speak; no completely different or even alien things, properties, events, or processes exist in them. So, what I said above is true of *I** as well; even relative to *I** causal connections cannot be invented by assuming new variables.

Secondly, possibility is here restricted by the given probability measure. The conception is not that each world carries with it its own probabilities (or its own laws of nature, etc.); worlds are taken here just as large conjunctions of singular facts. Hence, there is no way of inventing suitable causal connections by assuming probabilistic dependencies where there are none according to the given measure (or by referring to queer laws of nature, etc.).

This shows my variant of Reichenbach's principle to be nontrivial despite the fact that it considers potential relevance only. Still, I think that it would be more adequate and more interesting to refer to actual relevance. Hopefully closer variants can be proved as well on the basis of the above explication. For instance, suppose that *i* and *j* are two correlated variables, that ω is a world in which the actual circumstances of causal relationships are ideal in the sense explained in my (1990, sect. 4), and that *K* is the set of variables identical with or actually causally relevant to *i* in ω and *L* the set of variables identical with or actually causally relevant to *j* in ω. Then the independence $^{\omega}(K - L \perp {}^{\omega}(L - K) / {}^{\omega}(K \cap L)$ follows; indeed the proof is completely analogous to the one given above; the only difference is that the step from (6) to (7) now requires the additional assumption that circumstances are ideal in ω, an assumption which I have argued in (1990, sect. 6) to be important for characterizing causally well-behaved worlds. However, this result is not yet satisfying because that independence does not entail the desired independence $^{\omega}i \perp {}^{\omega}j / {}^{\omega}(K \cap L)$. Perhaps further assumptions of causal well-orderedness help. As I say, there is hope, and there is space to be explored.

Notes

* I am indebted to my commentator André Fuhrmann. The discussions with him led to a considerable improvement of the theorem in section 4, and his comments made clear to me in which respects this theorem may still be weak. Again, I am very grateful to Karel Lambert for rich advice in grammar, style, and content.

See, e.g., the variety of opinions presented in Bennett (1988).

[1] This characterization, or at least its non-modal part, is Quine's, reinforced in Quine (1985, 167). In (1969), Davidson thought it to be the second best view; and according to his assent to Quine in (1985, 175) it even is his present favorite.

[2] Of course, there are ways to get around this difficulty. For instance, one may associate with each event the proposition that that event exists or occurs, as does Lewis (1973). But even then it is doubtful whether the predominant view is appropriate. Lewis (1986, 241-269), in any case, prefers to say, at least for the present purpose, that an event is whatever engages in reasonable causal relations according to our best theory of causation.

[3] However, Kim's events rather behave like facts, as is critically remarked by Bennett (1988, ch. V).

[4] Indeed, in (1991, sect. 3) I have explained why I take the subjective interpretation as primary even in dealing with causation, and my (1992) is an attempt to do justice to the tendency to conceive causation objectively.

[5] In (1956, 159) he calls (2) and (3) assumptions. But later on on that page he says that "when we say that the common cause *C* explains the frequent coincidence, we refer ... also to the fact that relative to the cause *C* the events *A* and *B* are mutually independent." Hence, the conditional independence is indeed essential for *C* to explain the correlation and thus presumably also for *C* to be a cause, since causes are meant to explain.

[6] One exception is, of course, Salmon who has insisted, in (1970) and later writings, that within a probabilistic context negative relevance is explanative as well. I think he is right, but I also think there is not really a conflict. The reason is simply this: Within a deterministic context there seem to be only causes; counter-causes are quite strange or even impossible (though in my view the problem is not their impossibility, but their non-objectifiability in the sense explained in my (1992).) The case is different, however, with a probabilistic setting. Indeed, there it is not so clear how to disentangle causes and counter-causes; causal relevance as such, be it positive, negative, or mixed, is the much clearer notion. If others take positive relevance as a mark of causation, then negative probabilistic relevance certainly expresses causal relevance as well. And as Salmon insists, any explanation must attempt to state the causally relevant factors as completely as possible.

[7] Cf. van Fraassen (1980, 30); the point is repeated in his (1982b, 205). Indeed, the point is intended as a criticism of Salmon's interactive forks in which a residual positive correlation conditional on the common cause is assumed. In conversation Salmon has admitted that this remains a poorly understood, if controversial point in his discussions with van Fraassen.

[8] Of course, provability within a classical framework is not the same as analyticity. A nice example of this attitude is Davidson's individuation of events in (1969) which surprisingly entails a kind of causality prinicple, namely that each event with the exception of at most one has a cause or an effect.

[9] Cf. my (1991, sect. 4). The two principles do not have the same status in my view because the common cause principle starts from an assumption which will allow for a derivation of its conclusion, whereas the principle of causality makes an unconditional existence claim.

[10] Cf., e.g., Putnam (1987, 3-8), and Ben-Menahem (1990).

[11] For more details cf., e.g., Bell (1981), van Fraassen (1982a); and Skyrms (1984).

[12] My (1992) is an attempt to substantiate this hope within my framework.

[13] For proof see Dawid (1979), Spohn (1980), or Pearl (1988), sect. 3.1.

[14] The theorem may also be proved within the theory of Bayesian networks with the help the so-called criterion of *d*-separation and its relation to probabilistic independence; cf. Pearl (1988, 116ff.). This is so because $K - L$ and $L - K$ are clearly *d*-separated by $K \cap L$.

[15] At that time, "cause" often meant "total cause", because "total cause" was taken as the primary object of analysis and "partial cause" was hoped to be an easily derivable notion. This hope was not realized. So, nowadays "partial cause" has moved to the center of the analytic attempts, and "total cause" is the derivative notion. In the above explication "cause" clearly means only "partial cause".

[16] Salmon (1970, 220ff.) does so as well when he conceives explanation as the specification of causally relevant factors or of an explanatory partition. There, Salmon also requires that the members of an explanatory partition be homogeneous with respect to the explanandum. This has a counterpart in the present theory, namely in the fact that a variable is independent from its entire past given the set of all variables potentially directly causally relevant to it (this is, roughly, expressed in line (9) of the above proof); this set is maximally specific in Hempel's terminology.

[17] In a way it's worse. I defined direct potential causal relevance to be direct causal relevance in some world. Thus the transitive closure of the first is even weaker than the transitive closure of the second in some world; transitive closure and existential quantification do no commute.

[18] Cf., however, my (1991, sect. 4).

References

Bell, J.S. (1981), "Bertelmann's Socks and the Nature of Reality", *Journal de Physique 42*: 41-61

Ben-Menahem, Y. (1990), "The Inference to the Best Explanation", *Erkenntnis 33*: 319-344

Bennett, J. (1988), *Events and Their Names*, Oxford: Clarendon Press

Cartwright, N. (1979), "Causal Laws and Effective Strategies", *Noûs 13*: 419 -437; also in N. Cartwright, *How the Laws of Physics Lie*, Oxford 1983: Clarendon Press, pp. 21-43

Davidson, D. (1969), "The Individuation of Events", in N. Rescher (ed.), *Essays in Honor of Carl G. Hempel*, Dordrecht: Reidel, pp. 216-234

Davidson, D. (1985), "Reply to Quine on Events", in LePore, McLaughlin (1985), pp. 172-176

Dawid, A.P. (1979), "Conditional Independence in Statistical Theory", *Journal of the Royal Statistical Society, series B, 41*: 1-31

Eells, E., and Sober, E. (1983), "Probabilistic Causality and the Question of Transitivity", *Philosophy of Science 50*: 35-57

van Fraassen, B.C. (1980), *The Scientific Image*, Oxford: Oxford University Press

van Fraassen, B.C. (1982a), "The Charybdis of Realism: Epistemological Implications of Bell's Inequality", *Synthese 52*: 25-38

van Fraassen, B.C. (1982b), "Rational Belief and the Common Cause Principle", in R. McLaughlin (ed.), *What? Where? When? Why?*, Dordrecht: Reidel, pp. 193-209

Glymour, C., Scheines, R., Spirtes, P., and Kelly, K. (1987), *Discovering Causal Structure. Artificial Intelligence, Philosophy of Science, and Statistical Modeling*, San Diego: Academic Press

Good, I.J. (1961-63), "A Causal Calculus", *British Journal for the Philosophy of Science 11*: 305-318, *12*: 43-51, and *13*: 88; also in I.J. Good, *Good Thinking: The Foundations of Probability and Its Applications*, Minneapolis 1983: University of Minnesota Press

Granger, C.W.J. (1980), "Testing for Causality. A Personal Viewpoint", *Journal of Economic Dynamics and Control 2*: 329-352

Kiiveri, H., Speed, T.P., and Carlin, J.B. (1984), "Recursive Causal Models", *Journal of the Australian Mathematical Society, Series A, 36* : 30-52

Kim, J. (1973), "Causation, Nomic Subsumption, and the Concept of an Event", *Journal of Philosophy 70* : 217-236

LePore, E., and McLaughlin, B. (eds.) (1985), *Actions and Events: Perspectives on the Philosophy of Donald Davidson*, Oxford: Blackwell

Lewis, D. (1973), "Causation", *Journal of Philosophy 70*: 556-557; also in Lewis (1986), pp. 159-172

Lewis, D. (1980), "A Subjectivist's Guide to Objective Chance", in R.C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, vol. II, Los Angeles: University of California Press, pp. 263-293; also in Lewis (1986), pp. 83-113

Lewis, D. (1986), *Philosophical Papers, vol. II*, Oxford: Oxford University Press

Pearl, J. (1988), *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, San

Mateo: Morgan Kaufmann

Putnam, H. (1987), *The Many Faces of Realism*, La Salle, Ill.: Open Court

Quine, W.V.O. (1985), "Events and Reification", in LePore, McLaughlin (1985), pp. 162-171

Reichenbach, H. (1956), *The Direction of Time*, Los Angeles: The University of California Press

Salmon, W.C. (1970), "Statistical Explanation", in R.G. Colodny (ed.), *The Nature and Function of Scientific Theories*, Pittsburgh: The University of Pittsburgh Press, pp. 173-231

Salmon, W.C. (1978), "Why Ask, 'Why?'? An Inquiry Concerning Scientific Explanation", *Proceedings and Addresses of the American Philosophical Association 51*: 683-705

Salmon, W.C. (1980), "Probabilistic Causality", *Pacific Philosophical Quarterly 61*: 50-74

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

Skyrms, B. (1984), "EPR: Lessons for Metaphysics", *Midwest Studies in Philosophy IX*: 245-255

Spohn, W. (1980), "Stochastic Independence, Causal Independence, and Shieldability", *Journal of Philosophical Logic 9*: 73-99

Spohn, W. (1983), "Deterministic and Probabilistic Reasons and Causes", in C.G. Hempel, H. Putnam, and W.K. Essler (eds.), *Methodology, Epistemology, and Philosophy of Science. Essays in Honour of Wolfgang Stegmüller on the Occasion of his 60th Birthday*, Dordrecht: Reidel, pp. 371-396

Spohn, W. (1990), "Direct and Indirect Causes", *Topoi 9*: 125-145

Spohn, W. (1991), A Reason for Explanation: Explanations Provide Stable Reasons", in W. Spohn, B.C. van Fraassen, and B. Skyrms (eds.), *Existence and Explanation*, Dordrecht: Kluwer, pp. 165-196

Spohn, W. (1992), "Causal Laws are Objectifications of Inductive Schemes", in J. Dubucs (ed.), Theory of Probability, Dordrecht: Kluwer

Suppes, P. (1970), *A Probabilistic Theory of Causality*, Amsterdam: North-Holland