Salmon, Statistics, and Backwards Causation

Author(s): David Papineau

Source: *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1978, Vol. 1978, Volume One: Contributed Papers (1978), pp. 302-313

Published by: The University of Chicago Press on behalf of the Philosophy of Science Association

Stable URL: https://www.jstor.org/stable/192646

Salmon, Statistics and Backwards Causation

David Papineau[1]

Macquarie University

1. Introduction.

Smokers are more likely to develop lung cancer than non-smokers. This statistical association gives us reason to suppose that smoking can *cause* lung cancer. But we would come to reject this supposition if we discovered that when family histories of lung cancer are taken into account smoking makes no further difference to the probability of lung cancer: that amongst those with a family history of lung cancer (and amongst those without) non-smokers are just as likely to develop lung cancer as smokers. Such a discovery would lead us to conclude that, despite initial appearances, smoking actually makes no causal contribution to lung cancer, but is statistically associated with it only because a genetic predisposition to lung cancer also somehow predisposes people to smoke.

What is the rationale for such inferences? *Why* does a statistical association between two events constitute a *prima facie* reason for supposing one to be the cause of the other? And why is this conclusion discredited when we discover that the statistical association "disappears" when we partition our reference class by "controlling" for some prior factor? Our intuitions here are so strong that these questions seem silly. But remember that causes are standardly supposed to *determine* their effects. How can statistical data show us what determines what? The answer is by no means immediately obvious.

In his recent [3] and [4] Wesley Salmon has developed an approach to explanation and causation which promises to solve this puzzle. Amongst other things Salmon denies that causation need be deterministic. He suggests that it is quite enough for one event to cause another that they simply be statistically associated and this association not be "screened-off" by any prior factors.[2]

In this paper I shall attempt to provide a different kind of solution, one which does not involve rejecting the traditional view of causation as essentially deterministic. In so doing I shall cast some general doubt on Salmon's approach.

PSA 1978, Volume 1, pp. 302-313

To explain the relevance of statistical data to causal conclusions I shall need to refine the orthodox neo-Humean regularity analysis of causation. At the end I shall argue that the refinement in question helps with the problem of "backwards causation"; that is, that it gives us a time-independent analysis of causation which explicates the asymmetry of causal relationships with respect to temporal direction.

2. Barometers and Eclipsing.

The current neo-Humean orthodoxy is that one event is a cause of another if the former is one of a set of factors which is jointly and minimally sufficient for the latter. (For instance, see Mackie [1].) But then there is the problem of the barometer. Whenever the barometer falls, and it hasn't been fiddled with, and there is a certain amount of water vapour in the atmosphere, and various conditions for the formation of water droplets are satisfied, it rains shortly afterwards. The barometer's fall is part of a perfectly good minimally sufficient condition for rain. Yet the barometer's fall does not <u>cause</u> rain. On the contrary, our intuitions are that the barometer's fall and the rain are both independent effects of an earlier factor, namely a fall in atmospheric pressure.

This illustrates a general difficulty. Given the neo-Humean analysis of cause, if A has two successive effects B and C then it will standardly follow that B is a "cause" of C.[3] But we do not standardly want to allow this. To be sure, B <u>can</u> be a cause of C in such cases - A can cause C <u>by</u> causing B. Apart from affecting the barometer, the fall in pressure also causes a condensation of water vapour before the rainfall. And this <u>is</u> an intermediate cause of the later effect, the rainfall. But the difficulty remains. The neo-Humean analysis lacks a way of distinguishing these cases, where an initial cause produces some later effect <u>by</u> causing some intermediate effect,[4] from those where an initial cause produces two effects independently.

Our intuitions are that the barometer's fall, unlike the condensation of water vapour, is in itself irrelevant to the subsequent rainfall. In the circumstances rain would have followed the drop in pressure even if the barometer hadn't fallen. By contrast the rain wouldn't have followed the drop in pressure if there hadn't been any condensation.

Can we find a basis for this distinction? Note that on all occasions where a barometer fall occurs together with the other factors (the barometer's not being manipulated, high humidity, etc.) which make up a guarantee of rain, a drop in atmospheric pressure will also have occurred and, together with some of the same factors (high humidity, etc.), will also have made up a sufficient condition for rain. And what is more this latter set of factors will on occasion occur, and suffice for rain, even when no barometer fall occurs - as when there are no barometers around.

This seems to be characteristic of cases where B and C are independent effects of A. In a sense A <u>eclipses</u> B as a cause of C. Whenever

B "guarantees" C, A is there to suffice for C too; and in similar circumstances A suffices for C even when B is absent. So I would suggest that for B to be a real cause of C, and not an irrelevant side-effect of some prior cause, not only does B need to be part of a minimally sufficient set for C, but in addition there should be no other factors which eclipse B as a cause of C: that is, there should be no other minimally sufficient set not containing B which

(i) is always present when the set containing B is, and

(ii) is sometimes present when the set containing B is absent.

(In this paragraph, and from now on, "A", "B", etc. are to be taken as standing for types of events rather than for individual events.)

Venn diagrammatically, to require that nothing eclipses B as a cause of C is to require that there is no A satisfying the diagram:

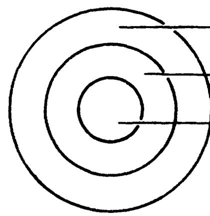

Consider now the case where B is indeed causally intermediate between A and C, Here B does seem to be crucially relevant to C's occurrence. A fall in pressure (A) suffices for rain (C) only insofar as it is accompanied by condensation (B); while condensation can suffice for rain even in the absence of a preceding drop in pressure, as when potential nuclei are artificially or naturally introduced into a supersaturated atmosphere. Indeed in this case it is B that eclipses A as a cause of C. Diagrammatically:

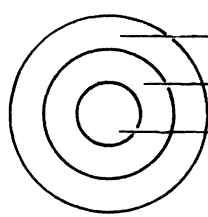

But does this mean that in such cases A is not a cause of C after all? This would be unsatisfactory. The drop in pressure surely causes the rain all right, albeit via causing condensation.

We can sort this out with the following definition. One event is a <u>direct cause</u> of another if, as above, it is part of an uneclipsed sufficient condition for it. One event is a <u>causal ancestor</u> of another if it is a direct cause or a direct cause of a causal ancestor. Then we can allow that causal ancestors other than direct causes are also causes, albeit indirect ones. So when A causes C through B, B will eclipse A as a direct cause of C - but nevertheless A will be a causal ancestor of

C, through causing B which causes C.

(One can imagine certain awkward cases where the no-eclipsing requirement fails to distinguish straightforwardly between intermediate causes and side-effects. These will be those where A and B always suffice for C together, where the Venn diagrams representing their actual occurrences coincide. For instance, suppose a certain virus can be followed by a skin infection and then by a fever, and suppose further that the virus is never found without the infection nor vice versa. Does the virus cause the fever by causing the infection, or does it cause them both independently? The obvious way to analyse this kind of case is by reference to a theoretical account of the mechanisms behind the co-occurrence of virus, infection and fever. This can be seen as a kind of elaboration of the no-eclipsing requirement. The theoretical explanation allows us to consider (theoretically deviant) contexts in which the virus occurs without the fever, and vice versa, and to use them to show whether the one would still suffice for the fever anyway, even with the other absent. Some further comments on this kind of case are made at the end of the Appendix.)

3. Causes from Statistics.

We are now in a position to understand how statistical associations can give us grounds for causal conclusions. Roughly the idea will be that a statistical association between Z and Y is an initial indication that Y is part of an uneclipsed sufficient condition for Z - but if this association "disappears" when we "control" for some V, then this means that V eclipses Y as a cause of Z.

Suppose[5] Y and Z are statistically associated. That is, Prob (Z/Y) > Prob (Z). The standard objection to concluding that Y is a cause of Z is that the association might be "spurious", a consequence of their both being independently associated with some prior cause. But suppose for a moment that Y is not itself associated with any other causes Z may have. It can then be proved that if Z is caused whenever it occurs Prob (Z/Y) > Prob (Z) if and only if Y is a direct cause of Z. A detailed proof is best left to the Appendix below. But the idea is simple enough. If Y is a direct cause, then when Y occurs we will get Z if either one of Z's other direct causes is present or if the factors which in conjunction with Y suffice for Z are present; while in the absence of Y we will get Z only if one of Z's other direct causes is present. So if Y is a direct cause of Z then Prob (Z/Y) > Prob (Z). Conversely, if Y is not a direct cause of Z then its occurrence will not increase Z's probability. But these inferences clearly require essentially that Y is not statistically associated with Z's other direct causes, for this could produce an association between Y and Z even if Y were not a cause of Z, and it might just lead to the absence of an association between Y and Z even when Y was a direct cause of Z.

What then if Y is associated with other direct causes of Z? To deal with this possibility we need to "control" for these other causes; that is, to consider separately classes of cases which are similar in respect of these other causes. Then, if Y still makes Z more likely _within_ any such classes, we can conclude that Y is indeed a direct cause of Z; whereas if it does not we can conclude that despite any original association Y is not after all a direct cause of Z. The detailed substantiation of this principle is once more best left to the Appendix. But again the idea is simple enough. We can go through the same reasoning as before, only now "within" those classes where we "hold constant" all those other causes that are confounding the association between Y and Z. By so doing we effectively ensure that Y is not associated with any "remaining" causes of Z, and so can be definite about what any remaining association between Y and Z shows about their direct causal relationship.

These results explicate our intuitive judgments about data like the hypothetical smoking-cancer statistics with which I introduced this paper. They show why Y's being statistically associated with Z is indeed a _prima facie_ reason for thinking Y is a cause of Z. At the same time, they show that this inference presupposes that Y is statistically independent of Z's other causes (or, to be more precise, that Y is statistically independent of all Z's causes other than those causal descendants of Y which are themselves in turn causes of Z - for if Y is an _indirect_ cause of Z it is only natural that its association with Z will be through its causal descendants). And so we also see now why we should conclude that Y isn't a cause of Z after all if and when we discover that their association disappears when we "control" for some further cause of Z (excepting, again, those further causes which are intermediate between Y and Z). Of course in practice causal conclusions of the kind in question will always be tentative, even when we have no doubts about the lawlikeness of our statistical associations. To be sure that Y _is_ a cause of Z we need to be sure we have taken into account all other confounding influences. Though we could have good grounds for thinking this, it will be difficult to achieve complete certainty. Less obviously, judging that Y is _not_ a cause of Z also strictly requires that we have taken into account all confounding influences - for a negative association between Y and some unknown further cause of Z might always just be such as to "cancel out" exactly Y's association with Z even when Y is a cause.

It is worth emphasising that the no-eclipsing addition to the regularity account of causation is essential to our explicating the principles of inference in question. This can be illustrated by considering the barometer example once more. Once we "control" for the prior change in the atmospheric pressure falling barometers don't further increase the chance of subsequent rain. But recall that a falling barometer can be part of a perfectly good minimally sufficient condition for rain. What it fails to be is an _uneclipsed_ part of such a set. This shows that it is only because causes are specifically uneclipsed parts of sufficient conditions that we can take the absence of a statistical association to show an absence of causation. (The Appendix shows exactly where the no-eclipsing clause is required in the proof that statistical independence implies absence of causation).

4. Salmon and Indeterministic Causation.

I mentioned in the introduction that Wesley Salmon has recently suggested that for one event to be the cause of another is simply for the former to be relevant to the latter's probability. Antecedent circumstances making an effect certain is for Salmon thus just the limiting case of causation. Of course Salmon has to avoid counting all statistical associations as causal connections. He doesn't want barometer readings causing rain simply on the grounds that they make it more likely. To deal with this he specifies that the explanatory factors we adduce should comprise a "homogeneous reference class" - the probability of the explanadum should not be alterable by including further factors in the explanans. So falling barometers are no good for causing rain, in that adding in prior drops in atmospheric pressure gives a better probability for rain; but drops in atmospheric pressure are not ruled out, since adding barometer readings to them makes no difference to the probability of rain. In such a case Salmon says the atmospheric pressure "screens off" the barometer as a cause of rain.[6]

Clearly there are similarities between Salmon's "screening off" and my "eclipsing". But there remain fundamental differences between his approach and mine. I have been specifically concerned with inferring causes from probabilities when the effect in question is always determined. But for Salmon we will not have a "homogeneous reference class" in such cases until we have included all determining factors. Thus on Salmon's account we are only entitled to infer causes from probabilities when the statistical relationship between cause and effect reflects an actual indeterminacy in the effect's occurrence.[7]

What are we to make of this? We have two superficially similar systems of procedures for selecting legitimate explanations, both hinging on what happens to statistical associations when we "control" for prior factors, but applicable to mutually exclusive sets of situations. I think this should make us suspicious of Salmon's approach. An initial difficulty for his line is that the kind of procedures in question do have an obvious plausibility in cases where our statistical analysis patently does reflect an incomplete knowledge of relevant causal factors, as is testified by much medical and social scientific research. Moreover, I have suggested a justification for using those procedures in such cases. Salmon defends his interpretation of the significance of statistical data largely by reference to illustrative examples. It is very tempting to conclude that his interpretation is plausible only because of its superficial similarity to the one I have defended in this paper, that he is trading on the implicit fact that his procedures would be valid if the statistical associations involved were, contrary to his specifications, only partial reflections of underlying deterministic relationships.

What of events that really are undetermined? I think that the right view of such events is that they simply are not caused, and that as such their occurrence is not explainable. To be sure, we can consider in such cases antecedent conditions which affect the chance of the event's occurrence. Undetermined events can be taken to have objective non-unitary chances of occurrence in the single case, in addition to the frequency-in-a-reference-class probabilities we can attach to determined events in the way I have been presupposing so far. And so we can treat the occurrence of such single-case chances as events in themselves, and ask what caused them. Indeed the discussion of causation in this paper will apply quite straightforwardly to such questions. But why an undetermined event itself occurred, given that it had the chance it did, remains a further question, and one that lacks an answer.

5. Backwards Causation.

The no-eclipsing elaboration of the neo-Humean analysis of causation allows the possibility that effects might precede their causes in time. I shall now attempt to show that it is nevertheless what we need in order to grasp why this possibility is never actualised.

Suppose two determined events, A and B, are statistically associated. This kind of dependency calls for explanation. Single coincidences are only to be expected. But if A and B turn out in general to occur together more often than is indicated by their separate probabilities of occurrence we want some explanation of what is going on.

This demand will be satisfied if we can show there is a "causal net" connecting A and B, where the strands of the net are relationships of direct causation showing either that A is a causal ancestor of B or that B is a causal ancestor of A or that they have a common causal ancestor, and where the "exogeneous" direct causes of the final and the intermediate effects in the net are all singly and jointly statistically independent. For given such a net, the kind of reasoning employed in Section 3 can be used easily enough to account precisely for the difference between Prob (B/A) and Prob (B) solely in terms of the explicit deterministic connections involved and the independent probabilities of the exogenous factors.

It seems to me reasonable to suppose that all statistical associations between determined events can be explained away by reference to such causal nets.[8] For, as was shown in Section 3., if two determined events, A and B, are statistically associated, then (including the special case where A is itself a direct cause of B) A must be statistically associated with B's direct causes. And, by the same argument, it must be associated with their direct causes. Moreover A will occur if and only if some of its direct causes occur. And they will occur if and only if some of their direct causes do. So A's and B's causal ancestors will be statistically associated backwards ad infinitum. Of course if A is a causal ancestor of B, or B an ancestor of A, or if they share a common ancestor, this will be perfectly understandable, for after some point they will actually share causal ancestors. But for A and B to be statistically associated without being causally connected in some such way would propagate an implausibly cosmic succession of general coincidences.[9]

Note now that the association between two events will "disappear" if we "cut" their link by "controlling" for any event intermediate in the net linking them; that is, if we partition our reference class into cases where that event occurs and cases where it does not. This is derivable from the earlier result that if one event is not a direct cause of another, then any association between them will disappear if we control for those direct causes of the second with which the first is associated. Thus in practice we find that the statistical association between falling barometers and subsequent rain disappears when we control for temporally prior drops in atmospheric pressure. Similarly the association between administering penicillin and the subsequent disappearance of the symptoms of an infection will disappear when we control for whether there is a temporally intervening elimination of the relevant germs.

It is however significant that, whenever we find that an association between some A and B disappears on controlling for some C, it is invariably the case that either C comes between A and B temporally or C precedes both A and B temporally - controlling for later events never removes associations between pairs of earlier events. Thus, for instance, the association between smoking and drinking does not disappear when we control for subsequent death by heart failure - even amongst those who do so die (and amongst those who do not) the smokers are more likely to have been drinkers than the non-smokers. Nor will the association disappear when we control for any other subsequent occurrence.

The significance of this is that it gives us reason to suppose that later events are never causes of earlier events; that is, that insofar as later events are parts of sufficient conditions for earlier events they must be eclipsed parts. For suppose that later events did cause earlier events. Then we would expect it to be the case that some of the causal nets accounting for statistical associations between pairs of factors had intermediate links which came after such pairs of factors in time, and consequently that some associations between pairs of earlier events disappeared when certain later events were controlled for. But we don't find this. So it seems that later events don't ever cause earlier events.

It is worth noting that, although this argument does not give any specific illustrations of the way later sufficient conditions are eclipsed as causes of earlier ones, it is not just an attempt to show that there is some "macro-statistical" temporal asymmetry in our universe. On the contrary, the argument takes a certain such asymmetry as a straightforwardly contingent and unexplained <u>premise</u> (earlier associations don't ever disappear when later events are controlled for) and invokes another "macroscopic" premise (any statistical association between events must be explicable in terms of a "causal net" connecting those events) in order to conclude that there must be temporal asymmetries in the specific deterministic causal relationships connecting particular events. The link between the premises and the conclusion is of course provided by the no-eclipsing account of causation which explains the connection between "macro-statistical" associations and causal relationships.

6. Appendix.

The first thing to be shown is that if Z is always caused and Y is statistically independent of any other direct causes Z may have, then Prob (Z/Y) > Prob (Z) if and only if Y is a direct cause of Z.

Let us call a conjunction of direct causes together making up a minimally sufficient condition for some effect a <u>full</u> direct cause. We can express the disjunction of Z's full direct causes in the form YW or X, where X is the disjunction of those full direct causes that do not contain Y as a part, and W the disjunction of those sets of conditions which together with Y give full direct causes of Z. Our question is now whether W has any members - <u>are</u> there any sets of conditions which together with Y makes up full direct causes of Z?

Since, by hypothesis, Z occurs if and only if one of its full direct causes does,

Prob(Z/Y) = Prob(YW or X/Y)
= Prob(YW or YX)/Prob(Y)
= [Prob(YW) + Prob(YX) - Prob(YWX)]/Prob(Y).

And so, since Y is independent of X,

Prob(Z/Y)= [Prob(YW) - Prob(YWX)]/Prob(Y) + Prob(X).

Now, Prob(Z) = Prob(YW or X)
= Prob(YW) - Prob(YWX) + Prob(X).

So, given that Y does not have a probability of 1, which would effectively mean it always happened and so was ineligible as a cause to start with, Prob(Z/Y) > Prob(Z) if and only if Prob(YW) > Prob(YWX).

If Prob(YW) > Prob(YWX), then Y must be a direct cause. For, if it were not, YW and <u>a fortiori</u> YWX would be impossible, and so both would have zero probabilities. So if Prob(Z/Y) > Prob(Z) then Y must be a direct cause of Z.

If, conversely, Prob(YW) ≤ Prob(YWX) - that is, if Prob(YW) = Prob(YWX) - then either Y is not a direct cause and both sides of the equation are zero, or X is certain once YW obtains. But if the latter were the case then X would eclipse YW as a cause anyway. For X and Y are statistically independent, and given that YW's being certain is ruled out to start with this means that X will sometimes occur without YW. So in either case, if Prob(Z/Y) is not greater than Prob(Z), then Y cannot be a direct cause of Z. (It is here that the no-eclipsing analysis of causation is crucial).

That was the first thing to be shown. I now want to consider the situation when Y <u>is</u> associated statistically with X. Let $V_1,\ldots,V_n$ be those parts of the disjuncts of X which are either singly or jointly associated with Y. Then T, the disjunction of sets of factors "left" after we "remove" $V_1,\ldots,V_n$ from X, will be statistically independent of

Y. What we now need to know is whether the presence of Y makes Z more likely <u>within</u> any category of instances which are all similar in respect of $V_1,\ldots,V_n$. That is, is Prob(Z/YV) > Prob(Z/V), for any V which is a combination of the presence and absence of the various $V_1,\ldots,V_n$?

If it is, then we can definitely conclude that Y is a direct cause of Z. For we can go through exactly the same reasoning as before within the class of instances given by our V. Except now we have assured ourselves that Y is statistically independent of the factors in T, the "other causes" of Z "within" our V. So if Y does make Z more likely it must be that, together with certain other factors, which might include elements of V, it makes up a full direct cause of Z.

If Y increases the chance of Z within some combinations of $V_1,\ldots,V_n$ but not within others, then this shows that the relevant elements of such combinations are themselves part of full direct causes containing Y. If, however, Y does not make Z more likely for <u>any such</u> combination - if Prob(Z/YV) = Prob(Z/V) for all possible V's - then we can in general conclude that Y is <u>not</u> a direct cause of Z. Again we can apply our earlier reasoning "within" each V with the assurance that Y is independent of Z's "other" causes: and so we can conclude that there is nothing which together with any combination of $V_1,\ldots,V_n$ can be added to Y to give a full cause of Z.

There is however one exception to the argument of this last paragraph which needs mentioning. Suppose Y <u>is</u> in fact a cause of Z, but that whenever it occurs it is made <u>certain</u> by a given combination of other causes V. Then there is no question of its increasing the probability of Z within that V. Technically Prob(Z/YV) = Prob(Z/V), since the condition YV will be effectively equivalent to the condition V. In the original simple case where we assumed that Y was independent of Z's other causes this difficulty did not arise because if Y were absolutely certain it would be ineligible as a cause. But now we have to worry about a certainty-relative-to-V, which cannot be so ruled out. This is in effect a version of the difficulty which arises with the kind of virus-infection-fever case mentioned at the end of Section 2. We can't tell directly whether or not Y (skin infection) is eclipsed by V (virus) because they always occur together. Correspondingly, in such a case it will not be possible to use statistical data to decide whether or not Y is a cause.

If causally related factors characteristically always occurred together in this manner, we would not have the kind of statistical data discussed in this paper, and the argument mounted in Section 5 above would be inapplicable to our world. But our world is not like that.

Notes

[1] I am grateful to Hugh Mellor and Graham Nerlich for comments on earlier versions of this paper. The argument about backwards causation was originally suggested by a conversation with Dan Gvirtzman. Further treatment of some of the issues herein is given in Chapter 3 of my forthcoming For Science in the Social Sciences, London: Macmillan, 1978.

[2] It should be noted that in [3] Salmon expresses some reservations about actually defining causation in terms of statistical relations alone (p.55, p.81). And in [4] he emphasises spatio-temporal continuity and the ability to "transmit marks" as further characteristic features of causal relationships. But the important point is that he allows that causation can be indeterministic. (See Suppes [7] and Sayre [6] for unequivocal defences of statistical definitions of causality.)

[3] See Mackie [2], pp.83-4, for a general demonstration of this point.

[4] The barometer problem is often mentioned in the literature. But so far as I know Sanford [5] is the only account of causation which suggests a satisfactory way of solving it. Though Sanford does not explicitly discuss the barometer problem, I take his account to be similar in spirit to the no-eclipsing account I develop here.

[5] I am interpreting "probability" as the limiting frequency of a kind of event in a hypothetically infinite reference class. I take such probabilities to be evidenced by the frequencies observed in random sampling from such reference classes.

[6] Salmon originally presented the screening off rule as independent of the homogeneity requirement. But, as he notes in his [3], p.105, the former is derivable from the latter.

[7] For Salmon's (admittedly hedged) preference for actual indeterminism over epistemic uncertainty see his [3], pp.79-80.

[8] Salmon similarly suggests that all statistical associations should be explained by causal relationships ("the principle of the common cause",[3], p.70, [4], p.124), and also that this means there is a temporal asymmetry to causation ([3], pp.65-76). But our different views of causation lead us to give a different basis and significance to these ideas.

[9] For simplicity I assume in this section that we are considering a deterministic segment of reality. But if we were not, the argument could still be made to apply by "replacing" any undetermined events by their chances of occurrence.

References

[1] Mackie, J.L. "Causes and Conditions." American Philosophical Quarterly 2(1965): 245-55 and 261-4.

[2] -----------. The Cement of the Universe. Oxford: Oxford University Press, 1974.

[3] Salmon, W.C. "Statistical Explanation." In Nature and Function of Scientific Theories. (University of Pittsburgh Series in Philosophy of Science, Volume 4.) Edited by Robert G. Colodny. Pittsburgh: University of Pittsburgh Press, 1970. Pages 173-231. (As reprinted in Salmon, W.C. et. al. Statistical Explanation and Statistical Relevance. Pittsburgh: University of Pittsburgh Press, 1971. Pages 29-87.)

[4] -----------. "Theoretical Explanation." In Explanation. Edited by S. Körner. Oxford: Basil Blackwell, 1975. Pages 118-145 with discussion to p. 184.

[5] Sanford, D.H. "The Direction of Causation and the Direction of Conditionship." Journal of Philosophy 73(1976): 193-207.

[6] Sayre, K.M. "Statistical Models of Causal Relations." Philosophy of Science 44(1977): 203-14.

[7] Suppes, P. A Probabilistic Theory of Causality (Acta Philosophica Fennica XXIV.) Amsterdam: North Holland, 1970.