Aleatory Explanations Expanded[1]

Paul Humphreys

The University of Virginia

1. Introduction

In my paper "Aleatory Explanations" (Humphreys 1981a) I presented the rudiments of a theory of explanation which had the following features:

1) Explanations of indeterministic phenomena are given by means of a specification of the contributing and counteracting causes involved;
2) Probability values are not themselves explanatory and explanations need not specify probability values;
3) Maximal specificity criteria are not part of the logical criteria of adequacy for explanations;
4) Partial explanations which are nevertheless true can thus be given;
5) Explanations of this kind show that the traditional 'X because Y' pattern of explanation is inadequate for indeterministic causal explanations and must be replaced by 'X because Y, despite Z'.

In that paper, I restricted myself to explanations of events and their complements, involving only direct causes in two-state Markov chains. Some of the arguments presented by Joseph Hanna in his contribution to this session (Hanna 1983) show, I believe conclusively, that to generalize the model to cover larger sets of possible outcomes, including the case of continuous valued variables, requires a reassessment of the motivation behind probabilistic causality. Although the original features of aleatory explanations remain untouched by this change the generalization given here is substantially strengthened by the revisions, and indicates how such explanations fit into a traditional pattern of explaining deviations from the natural state of affairs.

2. The Issue

Let us take the simplest theory of probabilistic causality available - one in which probabilistic relevance is a necessary condition for causal relevance, and in which the ontology is one of events. It is a qualitative theory, in that no measure of the degree of relevance is provided, merely whether one event C is causally relevant to another, E. This is the foundation of theories of causality given in Reichenbach (1956), Suppes (1970), Salmon (1980), Humphreys (1980), and Fetzer (1981) among others.[2] An important distinction is often made between positive probabilistic relevance and negative probabilistic relevance, and between their causal counterparts, contributing and counteracting causes. Informally, a cause is positively relevant if it increases the probability of the effect's occurrence, and negatively relevant if it decreases that probability. In everyday talk of causation we frequently make such a distinction. A rattlesnake bite is positively relevant to an individual's death, whereas administration of an anti-venom serum is negatively relevant to that death. In order for such causes to form the basis of adequate explanations, causal claims must be based on objective features of the situation. In particular, the distinction between contributing and counteracting causes must not be arbitrary. However, most of the theories of probabilistic causality cited above are based on a binary event ontology, and because of this a source of potential ambiguity in the classification of contributing and counteracting causes has been overlooked.[3] Unless this difficulty can be avoided, I believe that all theories of causation which rely on relevance relations are seriously flawed, and any theories of explanation based upon relevance relations, including those of Professors Salmon and Fetzer, have their objectivity threatened. I shall first present the problem which I shall call 'the problem of the potential ambiguity of the contrast case' in as strong a form as possible before attempting a solution.[4]

3. The Problem

The natural question to ask of someone who claims that positively relevant factors raise the probability is "Raise the probability relative to what?". Two suggestions have been made for the binary case. We can require that $P(E/C) > P(E)$ or that $P(E/C) > P(E/\bar{C})$.[5] Where only events and their complements are considered, these conditions are formally equivalent. This equivalence immediately runs into difficulties once we move to a tripartite partition of conditioning events $[C_1, C_2, C_3]$. Then it is easy to show that

$$(1) \qquad P(E/C_1) - P(E) = \sum_{i=2,3} P(C_i)\,[P(E/C_1) - P(E/C_i)].$$

Suppose that we were interested in deciding whether C_1 was a contributing cause of E, according to the two comparisons suggested above.[6] Equation (1) clearly allows $P(E/C_1) > P(E)$, $P(E/C_1) > P(E/C_2)$ but $P(E/C_1) < P(E/C_3)$. Consider a concrete example. E is recovery from an illness, C_1 is administration of a moderate dose of a drug X, C_2 is administration of a placebo, and C_3 is administration of a heavy dose of X. Because the drug is effective but has unpleasant side effects, we perform an objectively random drawing to determine which of C_1, C_2, C_3 occurs, by means of a fair die. Hence each of the probabilities in (1) is objective. Let $P(E/C_1)=0.4$, $P(E/C_2)=0.2$, $P(E/C_3)=0.9$, and

PSA 1982, Volume 2, pp. 208-223

$P(C_i)=0.33$ for each i. Now suppose E occurs and was preceded by C_1. Is C_1 a contributing cause of E? Elementary calculations give $P(E)=0.5$, and hence according to Suppes' criterion, administering a moderate dose of X is a prima facie counteracting cause of recovery. This appears to be an incorrect assessment, because C_1 contains an effective agent for producing recovery and surely contributed to the latter event. Not as much as would have C_3, admittedly, but we are not interested in the magnitude of the contribution. More important, and this is in essence the point brought out by Hanna, is the fact that whereas in the binary situation the contrast case in the second kind of comparison is unambiguously given by the complement of the putative cause, as soon as we consider any more detailed analysis of possible causes, whether or not C_1 is a contributing cause will depend, within a generalized comparison of the second kind - $P(E/C_1)$ versus $P(E/C_i)$ - on what we choose as the contrast event. Relative to C_2 (the placebo), C_1 is a contributing cause, but relative to C_3 (a heavy dose of the drug), C_1 is a counteracting cause of recovery. Which is the correct comparison? Without additional constraints, probabilistic causality gives us no answer. The problem is whether such constraints can be objectively given.

4. The Source of the Problem

I have put the problem so far in terms of the contrast between positive and negative relevance. The issue runs deeper than this, for it affects relevance relations in general. A simple adaptation of the example above will show that if administration of another drug Y, equivalent in its effects to a moderate dose of X, is event C_4, we can have $P(E/C_1) - P(E/C_4)=0$ with the other inequalities remaining the same.[7] Then relative to C_2, C_3, C_4, event C_1 is, respectively, positively relevant, negatively relevant, and probabilistically irrelevant. Thus, even those who are concerned only with relevance <u>simpliciter</u> will suffer from the ambiguity of the contrast case. For example, any relative frequentist attempting to use relevance relations to arrive at a homogeneous reference class will require a solution to this problem. Also, any attempt to generalize Fetzer and Nute's probabilistic conditionals by using a finer-grained partition of quantitative properties will expose that theory of causation to similar ambiguities. And it goes without saying that the situation is even worse when we use quantitatively valued variables.

The problem, once realized, has a familiar air to it. We are comparing an actual probability with a counterfactual situation - what the probability would have been had some other event preceded E.[9] It is not the use of conditional probabilities but the comparison with an indeterminate non-occurrent situation which produces the ambiguity, and unless we can avoid the kind of pragmatic features which have been associated with the analysis of counterfactual and subjunctive conditionals, relevance relations cannot provide an objective basis for explanations.

The potential solutions which are open to us seem to fall into the

following categories:

1) To drop the idea of probabilistic relevance altogether, and to use only probability values in explanations.
2) To formulate a theory of probabilistic causality which avoids the problem by using only occurrent events or states of affairs.
3) To attempt to find an objective contrast case.

The first of these seems to me to be impossible without having previously resorted to either the second or the third. The reason for this lies in the fact that objective interpretations of probability, be they relative frequency or propensity interpretations, rely on relevance relations via randomness conditions, objective homogeneity conditions, or maximal specificity conditions, and as already noted, the problem of the ambiguity of the contrast case involves relevance relations in general. Thus, although there may be independent reasons for holding, as Salmon does in his recent work (e.g., (1981)), or Fetzer does in his causal-relevance theory (1981), that the attribution of a determinate probability lies at the heart of probabilistic explanation, this approach cannot be adopted prior to a solution of the ambiguity problem, at least for single-case explanations. (I shall argue later that such approaches are undesirable on other grounds.) The approach I wish to take is the third of those mentioned above, and I shall explain in Thesis 2 below why the second solution cannot be adopted. To introduce the approach I shall take, consider the following deterministic analogy. Within classical mechanics, there was a clear conception of what the behaviour of a body would be in the absence of any external forces, which was provided by Newton's First Law. This 'natural state' of a material object was of course different from the 'natural state' attributed by Aristotelian theory to sublunary elements not undergoing violent motion, or to the aetherial element undergoing natural motion. These respective natural states needed no explanation, whereas deviations from them did. Such theories often include an account of what a neutral state free from causal influences is like. Similarly, although far less clear, is the conception of rational behaviour in humans as the natural state. Deviations from this natural state are the subject of explanation, whereas rational behaviour requires no explanation.

The situation with aleatory explanations is more complex. As I have mentioned elsewhere (1981b), the essential feature which distinguishes probabilistic causes from sufficient causes and necessary causes is that multiple probabilistic causes are frequently present. This entails that the absence of a probabilistic cause does not usually reduce the probability of the outcome to zero nor that the neutral value of the relevant variable is that value which results in the lowest probability for the outcome, for negative values of a quantity may be a counteracting cause, where positive values are a contributing cause. Thus for the stochastic case we must consider a comparison with the state in which only the factor at hand is neutralized, in order that the effect of the other factors is taken into account and to ensure that if an interaction effect occurs with some other factor present, which for example converts what is ordinarily a contributing cause into a counteracting cause, this will manifest itself. I suggest, then, that in general, the component Y=y of the state < Y=y,Z > is a contributing cause of the system having the value X=x if

$$P(X{=}x/Y{=}y,Z) > P(X{=}x/Y{=}y_o,Z)$$

where y_o is the neutral state of the factor Y. This appears to capture the original intention of comparing the presence of a cause with its absence, and it is easily adapted to the case of events from the variables case given above. Some examples will illustrate its use.

Consider the case again of the three possible treatments--a heavy dose of X, a moderate dose of X, and a placebo. Let the neutral event of "doing nothing" be denoted by N. Then, if P(E/N) = 0.2, we have that C_1, administration of a moderate dose of the drug, is a contributing cause of E, the recovery. A similar approach will do for the example of a car skid. Let the neutral state be a speed of 0 m.p.h. Suppose further that the actual speed of the car was 40 m.p.h., that the posted safe speed for a dry road was 30 m.p.h., that at or below the posted speed the probability of a skid is zero, and that above 30 m.p.h. the probability is a monotonically increasing function of the speed. Hence, relative to the neutral state, the speed of 40 m.p.h. is a contributing cause of a skid. The speed might have been higher, say 47 m.p.h., and that too would have been a contributing cause had it occurred instead.[10] Indeed, any speed greater than 30 m.p.h. would be a contributing cause of a skid. Both of these examples are of course simplified versions of ordinary phenomena, and a parallel scientific example may thus be useful. Within rather wide limits, the probability of an electron being emitted from a metal by thermionic emission is a function of the absolute temperature.[11] The neutral state here will be at a temperature of absolute zero, at which temperature the probability of thermionic emission is zero. Furthermore, for a fixed temperature, the probability of emission is a function of the chemical potential for the metal involved. The neutral state here is when the potential of the electron gas external to the metal is in equilibrium with the internal chemical potential, and hence any increase or decrease in that external potential will be, respectively, a counteracting or contributing cause of an electron emission. (See, e.g., Knuth (1966), pp. 114-118, 213).

(Both of these last two examples involve a mixed ontology, where the effect is an event, and the cause a value of a variable, but in the general case, both cause and effect will be values of variables.)

The specification of a neutral state will generally depend upon the scientific theory being used. It will usually reduce to the question of whether an objective criterion for a zero value for a quantity can be given, but it is important to note that the theory dependence does not make the specification subjective or pragmatic in nature. I do not believe that a purely probabilistic criteria can be given, because using stochastic independence as a criterion of neutrality, as in P(Y/X.Z) = P(Y/Z), will lead to the kind of averaging problem noted in the discussion of Thesis 2 below.

5. The Dynamics of Causal Explanations

The use of the neutral state approach to probabilistic causality should, I hope, show that explanations specifying contributing and counteracting causes need not be ambiguous, nor subjective or pragmatic, and that what goes into the 'X because Y, despite Z' framework is not based on arbitrary decisions. I shall now show that the second, third, and fourth claims of aleatory explanations--briefly, that maximal specificity conditions are not necessary, and hence that partial, yet true explanations not citing probability values can be given--provide a significant advantage over other proposed models of probabilistic explanation. Because of space constraints, I shall argue here for five closely connected theses to establish this claim.

Thesis 1: As a general rule of explanation for specific outcomes, we explain the actual only in terms of the actual and not in terms of unactualized possibilities. Because of this, sentences concerning the occurrence of particular events play a much greater role in explanations of indeterministic phenomena than they do for deterministic phenomena.

One advantage which causally-oriented explanations have over logically-oriented explanations is that they force us to deal with the dynamic elements of processes leading to specific outcomes. The deterministic case is, in a sense, very simple in this respect, for a single prior determining state will, together with the relevant laws, suffice to explain the outcome, and remote causes constitute as good an explanation as proximate causes, perhaps even better, because given futuristic but not historical determinism, a remote cause will automatically give us a determination of proximate causes, although not vice versa. With indeterministic phenomena, specification of what specific states occur between an indirect cause and its effect is essential for an adequate explanation. To see this, consider the simplest case of the skidding car example. A car enters a curve at 40 m.p.h. (A_{40}); its brakes lock (B); it skids on a patch of ice (C); the car ends in the ditch (D). (See Figure 1 overleaf.)

What is the correct explanation of the event D? It seems quite clear that any explicit use as part of the explanation of events that failed to occur, such as $\bar{B}$ or $\bar{C}$ is as illegitimate as trying to explain $\bar{D}$, an event which also did not occur. The first prohibition in fact follows from the second because we often want in turn to explain things that appear in the explanans set, and allowing non-occurrents in that set would preclude us from giving any further explanation for those non-occurrents. The neutral state view of causation, of course, makes reference to non-occurrent states in order to classify causes, although those non-occurrent states do not appear in the explanation itself. Thesis 2 asserts that this is unavoidable.

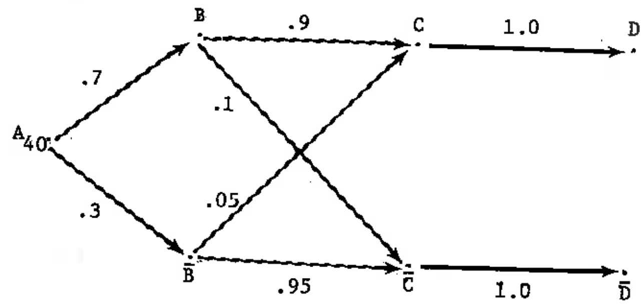

Figure 1

<u>Thesis 2</u>: In the decision regarding whether or not a state C is a cause of a state E, we cannot rely on an account of probabilistic causality which uses only occurrent states in the decision. Thus the second potential solution to the problem of the ambiguity of the contrast case is ruled out.

Suppose that we attempted to show that B was a cause of C in the car case by using only occurrent events. We try:

$$P(C/B.A_{40}) > P(C/A_{40})$$

Now, as well as the objection raised earlier that this kind of comparison led to wrong assessments in the drug example, we want to know what kind of probability $P(C/A_{40})$ is. It is in fact a weighted average:

$$P(C/A_{40}) = P(C/B.A_{40})\ P(B/A_{40}) + P(C/\bar{B}.A_{40})P(\bar{B}/A_{40}).$$

Thus we have in fact not adhered to our stated promise of using probabilities only of occurrent events; we need $P(\bar{B}/A_{40})$, for example. I believe that in any case where there are genuine alternative possibilities in the processes leading to the outcome, it is impossible to use only occurrent events in the probabilistic analysis of causes, and so our second kind of solution to the problem of the ambiguity of the contrast case is unavailable.

It might now be objected that a stronger version of Thesis 1 is required - that only the <u>closest</u> occurrent state is explanatorily relevant to the outcome state, especially when the process involved has the Markov property. If this were true, then serious problems would emerge, especially if a continuity principle for processes held for most phenomena. Once we see that a maximal specificity condition is not necessary for an adequate explanation, however, many of the reasons for insisting on a contiguity condition for explanations will disappear. First, I shall establish:

<u>Thesis 3</u>: <u>A principle of probabilistic continuity is not universally true</u> and even many macroscopic phenomena do not satisfy it.

Professor Hanna holds such a principle and has used it against Fetzer's account of explanation. Although various principles of continuity have a long tradition behind them, many microscopic processes do not seem to obey them.[12] Even in the macroscopic realm, they are often by no means obviously true. Consider the approach of the car towards the ice patch. Its path is continuous and indeed it is certain to hit the ice. However, it is not unreasonable to claim that at the point when the car skids, there is a probabilistic discontinuity. The shear forces and friction involved are such that at some point the bond between the tire and the ice breaks. Before the bond breaks, there is a probability of 0.9 that the car will land in the ditch; immediately afterwards there is a probability of 1.0. Another rather more idealized case of physical continuity not entailing probabilistic continuity is the classic example of the Landé blade.

Even without appeal to examples, principles of probabilistic continuity misrepresent, I believe, a fundamental feature of indeterminism. Consider a system in state S_i. It is in this same state from t_1 to t_2, and at t_2 it changes state discretely to S_k. Is this situation impossible? Russell (1912) claimed that it was in his infamous argument against the contiguity requirement in causation. Various principles of sufficient reason rule it out. But this case contains in a sense the essence of indeterminism. To deny it is to claim, as I believe Hanna does, and certainly Reichenbach did, that changes between states are always (at least macroscopically) made by continuous transitions. This entails that all macroscopic processes are continuous or "near" continuous. How plausible is this, and where does this idea come from?

There is no doubt that our conceptual categories are often discrete. Stocks are traded in discrete units, often of considerable size, divorces occur discretely, and tenure decisions often ride on small differences in the number of articles published. To argue that such discrete state models such as my having 1000 shares of IBM at noon, followed by my having none at 12:01 p.m. must be supplemented by a continuous set of underlying states is not only to commit oneself to a massive reductionist program but to conclude that explanations based upon such discrete states are to be excluded. What would be the basis of such an exclusion? It would seem that the reason would have to be that the theory involved was false, rather than merely incomplete. Incomplete explanations are not automatically false, of course. The explanandum sentence itself will often simplify the phenomena, as when we ask why Venus moves in (almost) a conic section, rather than why it moves in (almost) an ellipse. It is also not <u>false</u> to claim that crossing tall yellow peas with dwarf green peas explains why the first generation is composed exclusively of tall yellows, even though the explanation may well be considered incomplete. So why should incomplete probabilistic explanations be rejected?

6. Truth and Maximal Specificity

The reason is that:

traditionally truth in explanation entails maximal specificity,

hence omitting probabilistically relevant factors resulted in a false explanation. But this entailment only holds if we insist that an essential part of an adequate explanation is the specification of the correct probability attached to an event or state. Let us call a maximal specificity condition any condition which requires that all probabilistically relevant factors must be cited in an objective, non-epistemically-relativized explanation. Consider the statistical-relevance model, which Salmon has now abandoned, but is the clearest example of this. To explain an event under the S-R model it is necessary to assign it to the correct broadest objectively homogeneous reference class. Omitting a probabilistically relevant factor from the specification of the reference class would render it inhomogeneous. Hence a maximal specificity condition is required for S-R explanations of single-case events. Within Fetzer's single-case propensity theory of explanation a requirement of maximal specificity (and a requirement of strict maximal specificity) is imposed in order that the true law-like sentences in the explanans incorporate all and only nomically relevant properties. Hempel's inductive - statistical model also imposed a requirement of maximal specificity in order to avoid the problem of explanatory ambiguity.[13]

Within each of these models, the omission of even a single probabilistically relevant factor will result in a false probabilistic law being used in the explanation. This property applies also to the models given in Jeffrey (1969) and Railton (1978). This is not so in a model where explanations consist in a specification of the contributing and counteracting causes. Professor Fetzer has claimed in his paper (1983) that "the classification of contributing and counteracting causes, after all, is dependent upon the complete set of relevant properties present at that singular trial," and that "application [of Humphreys' canonical format] presupposes the satisfaction of other conditions whose fulfillment possesses greater theoretical and intuitive importance, namely: the Requirement of Maximal Specificity (for lawlike sentences that are true)...". (p. 203).

These claims are simply mistaken. Maximal specificity requirements can be imposed for either epistemological reasons or as truth conditions. In the first case, they are often used as an epistemological insurance policy, guaranteeing that our explanation will not be overthrown by the discovery of previously unknown relevant factors which alter the value of the probability in an undesirable way. This seems to have been Hempel's original reason for imposing it, because in his model we must be sure that there will be no discovery of factors which changed the inductive probability from a high to a low value. There seems to be no reason why the discovery of new factors which raise a high probability even further cannot be allowed, though, and a partial

maximal specificity condition seems possible on a Hempelian model.

The more important of the two reasons is the role played by maximal specificity as a truth condition. I believe that Professor Fetzer has failed to appreciate that his model of explanation cannot allow as an explanation any argument in which even a single relevant factor has not been cited; not because anything less would be incomplete, but because anything less would be false. Thus, if an individual dies from lung cancer, having been a heavy smoker, omitting from the explanation the following probabilistically relevant factors will result in a false probability value being given:

i) cosmic radiation from α - Centauri,

ii) an hereditary characteristic inherited from one's great-great-grandfather,

iii) particles from a smokestack in Salem, Oregon.

Of course as a practical matter, we omit these, but the important point is that aleatory explanations can omit such minor factors and remain true, whereas models citing probability values cannot. We thus arrive at:

Thesis 4: Explanatory models which use probability values as part of the explanation cannot distinguish between complete explanations and true explanations.

This in itself is an excellent reason for dropping the probability value requirement, and aleatory explanations avoid the problem because of the existence of partially or fully invariant contributing and counteracting causes. Whatever the conditions, increasing the frequency of incident radiation over the zero-point external field level increases the probability of photoelectric emission, and increasing the atomic weight over the neutral value of lead increases the probability of α -decay. Furthermore, unless causal factors did operate independently of one another, randomization techniques would be impossible to use and so, to a large extent, would experimental techniques where it was not possible to hold all the other variables fixed. It is not necessary to insist that all causes are fully invariant, as Mill did. As long as defeating conditions of a causal factor are absent (whether or not we know this) then irrespective of what we might later add to the explanation, the original explanation will stand as true. So, in particular, even if in many cases we are eventually able to supplement a discrete state model with an underlying continuous model, the original model can still have made a correct attribution of contributing and counteracting causes, and hence have given a true if incomplete explanation, when adding the continuous model does not convert contributing causes to counteracting, or vice versa. The probability values may even approach arbitrarily close to unity and the original causal claims remain true. Thus we have:

Thesis 5: Discrete state models of phenomena can serve as the basis of true aleatory explanations, even though there is reason to believe that continuous processes underlie the states used in the model.

I conclude with three observations. Aleatory explanations will usually involve a combination of deterministic and indeterministic processes. Thus, although for example, it is determined that the car will land in the ditch once it starts skidding, now we have seen that the contiguity condition can be avoided, we can still cite the brakes locking as a genuine contributing cause of the skid, and hence of the car's landing in the ditch. Secondly, probabilistic causality does use probability values in the comparison between the actual state and the neutral state. This does not lead us back to the necessity of a maximal specificity condition. For if inf $P(E/XZ) > \sup P(E/NZ)$ where the infimum and supremum are taken over the possible values of the other unspecified factors Z, then specification of, or knowledge of, the actual value of the probability is not necessary for a correct classification of X as a contributing cause. Finally, it is not true to assert, as Professor Fetzer does, that I am interested only in explanation sketches. Every true aleatory explanation is an explanation in its own right, but one which may be improved upon by citing other causes as they are discovered. The condition of maximal specificity which Fetzer and others have imposed upon truth explanations is so extreme that I strongly suspect that very few attempts at explanation on their grounds will ever succeed in meeting the standard.

7. Summary

The evolution of attempts to provide a satisfactory framework for explanations of non-deterministic phenomena began with their being called "statistical explanations". The incorporation of theoretical factors led to that overly empirical label being dropped in favour of "probabilistic explanations". The conclusion I wish to draw is that probabilities are a means and not an end to many explanations. Of course, I am not claiming that causal explanations are the only kind possible; that would be foolish. I do believe that it is time to recognize probability for what it is in explanatory contexts - a widely used tool, rather like Fourier series or differential equations, and that it has no more of a central role in explanations than do those techniques.

Notes

[1]I am indebted to James Fetzer and Joseph Hanna for many helpful comments and an admirably objective exchange of views. This paper was written while the author held a University of Virginia Summer Research Award.

[2]A survey of some of these theories is to be found in Salmon (1980). It is important to note that Salmon's most recent published work (1981) on explanation places a great deal of emphasis on causal interactions and processes capable of transmitting marks, with less emphasis on probabilistic relations. Thus, when I cite earlier papers by Salmon, they are to be taken as illustrative of one particular approach, rather than one currently held by him.

[3]Suppes' theory (1970) sometimes uses larger partitions of the outcome space, although as we shall see, it is open to other objections given below. Four non-equivalent definitions of causality covering quantitative variables are given in the penultimate chapter of his monograph. The definition of a quadrant cause is modelled on the event case, and is open to similar objections. Regress causality, positive likelihood ratio dependence, and functional causation seem to have a different motivation from the event case, and to be designed for causal generalizations rather than single case causation. There is reason to believe that this distinction between the event case and the variables case in Suppes' work mirrors an early dispute between G.U. Yule and Karl Pearson over the correct analysis of measures of association. I hope to pursue this issue in a future paper. Work by Esary, Proschan, and Walkup (1967) provides additional definitions of association between variables. Because of the large number of measures of association used in the statistical literature, the choice of which often depends upon the specifics of the distribution used, it seems advisable to limit a theory of causality to only those basic features shared by most such measures. Furthermore, it is worth noting that we cannot make distinctions in the case of binary valued variables which we can in the more general case. When we have indicator variables of events; positive covariation of the variables, positive covariation of the variables under non-decreasing transformations, "associated variables", positive quadrant dependence, and positive regression dependence are all equivalent, which is not true in the general case. Furthermore, it is easy to show that positive covariance is equivalent to positive prima facie causation in the binary case. The fact that we are unable to make distinctions in the simple event case might be seen as an indication that we should not use it as a guide to correctness in the more general situations. However, if we can provide a consistent extension of the application of the principle of relevant difference to the general case, the fact that our causal judgements are far firmer in these simple cases than they are in the more complex situations is a good reason for using them as a basis for extending the theory along the same lines.

[4]This problem was, to the best of my knowledge, first raised explicitly in a previous draft of Professor Hanna's (1983) contribution. Because in the published version, he has emphasized other objections to my account of explanation, I have thought it worthwhile to bring out this issue, which is of critical importance for causal theories of explanation. The particular formulation is, however, my own. There appears to be a realization of the ambiguity problem in f.n. 24 of Salmon (1980), although no discussion is given there.

[5]The first of these is adopted in Suppes (1970), Ch 2, and the second

in Humphreys (1981a), p. 229.

6 In what follows, everything applies mutatis mutandis to counteracting (negative) causes. The generalized form of (1) is $P(E/C_j) - P(E) = \sum_{i \neq j} P(C_i)[P(E/C_j) - P(E/C_i)]$. Hence the criteria are unambiguous only when C_j is a maximal contributing cause or a minimal counteracting cause.

7 Let $P(C_i) = 0.25$ and $P(E/C_4) = 0.4$. Then $P(E) = 0.475$.

8 See, e.g., Fetzer (1981), especially the disclaimer on pp. 51-52.

9 I assume here that the probability distribution is a deterministic function of parameters and (fixed) conditioning events. Some distributions, such as the compound Poisson, do not satisfy this assumption, but I shall ignore such distributions to avoid complicating the discussion, for they seem to require more than a straight "would" counterfactual in this sentence.

10 In a private communication, Professor Hanna suggested that we take as the contrast case the state immediately preceding the state we are considering as a candidate for positive relevance. He rejects it on grounds of lack of generality, but I believe it is unsatisfactory for different reasons. Suppose that immediately prior to entering the curve, the driver slowed from 50 m.p.h. to 40 m.p.h. Because there is reason to believe the situation is Markov, it is the fact that he hit the ice at 40 m.p.h. which is important, and what was the case beforehand is irrelevant. Two further points are worth mentioning here. We are using states rather than changes of state as causes; and it is not a change in speed which is causally relevant, but the actual speed. So the fact that he slowed is not a cause. Second, the fact that the driver slowed may be relevant to attributions of blame, but that is part of the explanation of how the state of 40 m.p.h. came about, rather than an explanation of how the skid occurred.

11 If N is the total number of free electrons in the metal, the probability of emission per unit time is

$$J'_N = \frac{4 \pi m}{h^3} (kT)^2 e^{-w/kT}$$

where w is the internal potential minus the zero-point energy. (See, e.g., Born (1962), Appendix XXXVI.) An important caveat must be made here. In many derivations of quantum statistical distributions such as the above, Bose-Einstein or Fermi-Dirac statistics are used, with the attendant assumption of particle indistinguishability. Application of probabilities to individual particles seems to be precluded by the indistinguishability, and so single case explanations will be difficult in these cases.

12 Leibniz, of course, held a continuity principle as a basic metaphysical truth, closely allied to his principle of sufficient reason.

Two of Russell's postulates of scientific inference, the postulate of quasi-permanence and the postulate of spatio-temporal continuity, are both continuity principles. (Russell (1948), pp. 506-510).

13 In fact, although I shall not argue it here, had the requirement of the truth of the probabilistic laws occurring in the explanans been rigorously applied, a maximal specificity requirement would have been a logical consequence of the truth of those laws.

References

Born, M. (1962). Atomic Physics. 7th ed. London: Blackie and Son.

Esary, J.D.; Proschan, F; Walkup, D.W. (1967). "Association of Random Variables, with Applications." Annals of Mathematical Statistics 38: 1466-1474.

Fetzer, J.H. (1981). Scientific Knowledge. Dordrecht: D. Reidel.

------------. (1983). "Probabilistic Explanations." In PSA 1982. Volume 2. Edited by P.D. Asquith and T. Nickles. East Lansing: Philosophy of Science Association. Pages 194-207.

Hanna, J. (1983). "Probabilistic Explanation and Probabilistic Causality." In PSA 1982. Volume 2. Edited by P.D. Asquith and T. Nickles. East Lansing: Philosophy of Science Association. Pages 181-193.

Humphreys, P.W. (1980). "Cutting the Causal Chain." Pacific Philosophical Quarterly 61: 305-314.

--------------. (1981a). "Aleatory Explanations." Synthese 48: 225-232.

--------------. (1981b). "Probabilistic Causality and Multiple Causation." In PSA 1980. Volume 2. Edited by P.D. Asquith and R.N. Giere. East Lansing: Philosophy of Science Association. Pages 25-37.

Jeffrey, R. (1969). "Statistical Explanation vs. Statistical Inference." In Essays in Honor of Carl Hempel. Edited by Nicholas Rescher. Dordrecht: D. Reidel. Pages 104-113.

Knuth, E.L. (1966). Introduction to Statistical Thermodynamics. New York: McGraw-Hill.

Railton, P. (1978). "A Deductive-Nomological Model of Probabilistic Explanation." Philosophy of Science 45: 206-226.

Reichenbach, H. (1956). The Direction of Time. Berkeley: University of California Press.

Russell, B. (1912). "On the Notion of Cause." Proceedings of the Aristotelian Society 13: 1-26.

-----------. (1948). Human Knowledge: Its Scope and Limits. London: Allen and Unwin.

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

------------. (1981). "Causality: Production and Propagation." In PSA 1980. Volume 2. Edited by P.D. Asquith and R.N. Giere. East Lansing: Philosophy of Science Association. Pages 49-69.

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