MARIAM THALOS

THE REDUCTION OF CAUSAL PROCESSES

ABSTRACT. The principle that causes always render their effects more likely is fundamental to the enterprise of reducing facts of causation to facts about (objective) chances. This reductionist enterprise faces famous difficulties in accommodating common-sense intuitions about causal *processes*, if it insists on cashing out causal processes in terms of streams of events in which every event that belongs to the stream is a cause of the adjoining event downstream of it. I shall propose modifications to this way of cashing out causal processes, still well within the reductionist faith. These modifications will allow the reductionist to handle processes successfully, on the assumption that the reductionist proposal is itself otherwise satisfactory. I shall then argue that the reductionist enterprise lies squarely behind the Theory of Relativity, and so has all the confirmatory weight of Relativity behind it. However this is not all good news for reductionists. For throughout I shall simply assume that the reductionist proposal, to the effect that causes are just chance-raisers, is correct. And I shall sidestep problems with that proposal as such. And so I shall show that, if in the end we find the reductionist proposal unsatisfactory, it cannot be on grounds of its treatment of causal processes as such. Thus, while I shall argue that causal processes pose no extra trouble for redutionists, I shall be making a case that all the action between reductionists and their opponents should be focused upon the proposal to reduce the two-term causal relation itself to relations amongst probabilities.

1. INTRODUCTION

A cause functions to render its effects more likely. This proposition lies at the center of the enterprise of reducing facts of causation, to facts about objective chances. Boosters of the enterprise (probabilists, as I shall call them) propose to equate causes with their functions – with what crassly we may call the 'summable' or 'quantitative' roles they play in bringing about the patterns of events we observe – and thus (equivalently, as they see it) with the regularities purely that are manifested in nature. It is a Humean approach to causes. (However it is possible to maintain, as Cartwright (1989) and Sober (1988) do, that a probabilistic theory of causation will not reduce causal relations to relations in terms of probabilities, but only illuminate the nature of the causal relations by showing how they are intertwined with probabilistic relations. But this is not the position I am here concerned with examining and advancing.) The Humean approach faces difficulties in accommodating common sense on the subject of causal

Synthese **131:** 99–128, 2002.

processes. I propose to show that probabilists are in a position to do no worse at accommodating causal processes than at accommodating simple two-term causal relations of the form: C causes E. In other words, I will for my purposes simply assume that the probabilist's account of *'C causes E'* is acceptable, and proceed to show that probabilists can handle causal processes in an equally acceptable fashion. I shall conclude that, opinions to the contrary notwithstanding, causal processes pose neither extra, nor special work for probabilists, over and above what already they have to do to satisfy us on the subject of 'C causes E'. And that therefore focus upon processes as such simply diverts attention from the question whether the reductionist account of 'C causes E' should be accepted.

My strategy will be to show that the kernel of the notion of causal process is already within reach of the primary instrument probabilists utilize to diagnose 'C causes E' – namely, the notion of a screener-off. I will argue that invoking this notion can help probabilists accommodate certain intuitions about causal processes that are thought to be problematic for them. And I also will argue that, in fact, the connection between screening off and causal processes is the *best* grounds one can have for being a probabilist. For it has all the weight of the theory of Relativity behind it.

2. THE PROBABILIST STRATEGY

Probabilists, like everybody else, have grown up with the knowledge that a statistical correlation between two factors, such as for example smoking and heart disease, does not all by itself imply causation between them. And probabilists, no less than their fiercest opponents, are typically realists about causation: they adhere to the proposition that causes are manifested in the world of everyday experience, and that the cause-effect relation exhibits the central features we common-sensically attribute to it. Probabilists will be last to side with the lot who wish to do away with the notion of causation, as philosophically disreputable. And for probabilists, just as for their opponents, doing justice to causal relations between factors or events has something to with giving an account of arrow-like somethings (of possibly different shapes and sizes) between these factors or events. Probabilists depart from their opponents, however, in maintaining that we can achieve a complete understanding, as well as a full representation of these arrow-like dependence relations, by confining attentions exclusively to numerical probabilities, that can in turn be worked out purely from relative frequencies. And all this is because they don't believe there's more to causal reality than a certain structure or configuration of objective chances, which is not necessarily immediately evident from limited information about correlations but which can be got in due time with patience, luck and self-application.

How do probabilists say the reduction works? For the sake of concreteness I will adopt David Papineau's approach in (1989, 1993), which is perhaps most transparent for the reductionist strategy. Other probabilist schemes differ in matters of detail or emphasis, but all remain true to the basic principle: causes raise the probabilities of their effects, and we can measure that incremental enhancement of an effect's chances of occurring, which is due to a certain one of its putative causes, by measuring the effect's elevated frequency in populations where (very roughly) everything but that putative cause is held fixed. Here now is (more or less) Papineau's formula.

First, we need to define the notion of screening off:

DEFINITION 2.1. S screens off E from C if $P(E|C\&S) = P(E|S)$.[1]

And now we use the following three axioms, to be invoked as they apply, in the process of inferring causal relations from relative frequencies:

AXIOM 1 (A1). Pairs of events which are unconnected (either directly or through a common causal ancestor) will be uncorrelated: they will cooccur at chance; so $P(A\&B) = P(A) \cdot P(B)$.

AXIOM 2 (A2). Pairs of events which either have a common causal ancestor, or are such that one causes the other, are (at least partially) correlated: they will cooccur at better than chance: $P(A\&B) > P(A) \cdot P(B)$.

AXIOM 3 (A3). Causal ancestors and causal intermediaries screen off the correlations we expect to see by A2; so that if S is either a common cause of A and B, or a causal intermediary between them (for example, if A causes B *through* causing S which causes B), then it will be true that $P(A\&B) > P(A) \cdot P(B)$, but also that $P(A\&B|S) = P(A|S) \cdot P(B|S)$.

The basic idea, on this formula, is that of screening off, in terms of which the three axioms are intended to do their work. The notion of screening off is itself couched in probabilistic terms – and thus reducible to a structure of objective chances. A1 and A2, together, tell us that where there are no causes at work, there are no correlations, and also no screeners-off. And then A3 – the primary reductionist insight – tells us how that we must work backwards from the identification of screeners-off, to the identification of common causes and causal intermediaries. A3 is put purely in terms of screeners off (everything that comes after the semi-colon is purely by way of explicating everything that comes before). This fact – that the central insight of probabilism can be captured in terms of a single fundamental notion, that can itself be spelled out purely in quantitative terms – is quite important, as we will discover. For it is the ground of the argument I shall make, to the effect that the quantitative account of causality can be transformed with no remainders into a qualitative account.

This particular formula, in terms of a procedure of working backwards from correlations to causes – one that can be carried out recursively – also makes it absolutely transparent that there will be trouble. A3 reveals that we might run into difficulties when it comes to discerning, when we have a screener-off, whether it is a causal intermediary within a causal chain or one of a number of joint effects of a common cause. And we will have similar troubles when there are different and independent causal routes to a particular event.Thus the philosophical task reductionists undertake, for purposes of illuminating the precise and detailed causal relations between any two events, is to work out strategies for discriminating common causes, from intermediaries, from cases involving alternative routes. However, philosophical reductionists are not working in isolation on this task. They are not the only ones concerned with making these discriminations. There is an entire academic field, within mathematical statistics, devoted to working out strategies for making the discriminations reductionists want made, by cleverly designing experiments to produce the right correlational data for comparisons. The point of designing such experiments is to identify more and more causal detail, about intermediaries, alternative routes, and the like, so that the inference from frequency data to causes becomes, ultimately, indisputable. The strategies are discussed by Spirtes et al. (1992), and lightly sketched also by Papineau (1993).

For good reasons, the strategies and results of the discipline of statistics enjoy wide application across the sciences. We will have occasion to mention this work again, with good effect. But the point for now is simply this: those who undertake to identify strategies for making discriminations among for example common causes and intermediaries, undertake it for the purpose of illuminating the two-term causal relations of the form ‘C causes E’ or ‘C is *a* cause of E’. I would like simply to suppose that their strategies and principles for inferring causal relations from relative frequencies are acceptable. I want to show that this is similarly acceptable for the purposes of reducing causal processes as well, in the spirit of the probabilist proposal. And, in fact, I wish to show that the Theory of Relativity itself stands behind this reductionist strategy. But the first order of business is to rehearse the famous problems with this proposal.

3. TROUBLE WITH PROCESSES

The difficulties probabilists face vis-à-vis processes, are caused by examples that are carefully crafted variations on problematic cases to which Suppes (1970) called attention when formulating the statistical relevance (SR) account of causality – the ancestor of contemporary probabilistic accounts. These examples involve so-called *negative causes*.[2] In one of these variants – not the best, for reasons that will become clear, but one that will convey nicely the sense of the problems for now – a golfer hits a ball towards a cup (see Sober (1988); Eells and Sober (1983)). The shot is a good one, very likely to carry the ball to its mark. But a pesky squirrel runs out into the path of the ball and kicks it as it goes past. The kick – as we discover upon careful analysis of the motions of ball and foot, rendered on film – makes it less likely than before that the ball will achieve its mark. Even so, the ball comes through for the golfer. Does the kick help *cause* the ball to do so? What are the consequences, for our treatment of causal processes, of holding that the squirrel kick is *not* one of the causes of the ball's serendipitous arrival?

Everyone is agreed (let's suppose) that the squirrel kick makes a negative contribution to the golf ball's *chances* of reaching the cup, in our sample example.[3] What is in dispute is whether the kick is also a *cause* of the ball's making its mark. If we say it is, then we are entitled, or so it would seem, to the proposition that the ball enters into a *single, unbroken, causal process* from tee to cup. If, on the other hand, we say that lowering the probability of the ball's reaching its destination, counts against something's being a cause, then we seem to lose title to that proposition. For it seems we shall be compelled to say that the causal process inaugurated at the tee is broken when ball encounters squirrel, and possibly a new (causal) process begun then. The question, then, is whether the occurrence of a probability-lowering event terminates one process and inaugurates another. (Notice that this question does not arise when the event whose probability is lowered *fails to occur:* that same squirrel kick, if never followed by a drop of golf ball into cup, will not prompt the judgment that a process breakage has taken place.)

If, on the one hand, we are reluctant to go along with the proposal that the squirrel kick terminates the causal process begun at the tee, it will be because we are prepared to say that the kick is a true cause, but nonetheless contributes *negatively* towards production of the effect. Thus that the squirrel kick is a *negative cause* of the ball's dropping into the cup. If we embrace negative causes, however, we will not be in a position to call on facts about objective chances alone to make a discrimination between the causal process and the noncausal, for we will in effect be embracing the idea that events in a causal process do not have to make a positive incremental contribution to the occurrence of succeeding events, in order properly to belong to it. And so we shall have to call on some *independent* notion, such as for example Salmon's mark criterion,[4] or something more elaborate like the transference theory or the conserved quantity theory,[5] to discriminate the causal process from the noncausal. Or we will simply, and even less satisfactorily, have to take the notion of causal process as a primitive one, *sui generis*, and subsequently rely on intuitions alone to make the discriminations. And of course we shall then have to abandon the enterprise of reducing the facts of causation to facts about objective chances.

On the other hand, we might find instead that our sympathies lie with a certain style of philosophical system rather than with intuitions about ball-and-squirrel examples, and that we are more prepared to put faith in the guidance that objective probabilities provide, than in whatever initial reactions are had, universally or not, to ball-and-squirrel stories. In that case we will be hostile to the very idea of a negative cause, as a kind of oxymoron. In addition we might also prefer to retain the ability to discriminate the causal process from the noncausal purely on grounds that all events in a causal process contribute positively towards bringing about successor events. For we might be inclined to say that all events in a causal process must contribute positively towards production of an effect that in turn contributes positively to its successors; otherwise, what could their claim to belonging to a causal process rest on? And thus we would be in a position to make the discrimination between causal and noncausal process rest on the probabilities themselves, however they might shake out, and not have to rely on the services of such squishy concepts as marks or transfers for making the discrimination. And we might prefer to preserve also the possibility of reducing causes to objective chances. But as Papineau (1989, 344–46), who recommends this approach, points out, we shall also have to confess that there are fewer causal processes than initially supposed. Or, to put it more precisely, we shall have to confess that the causal processes we used to acknowledge as single and of an appreciable duration, are at best artificial concatenations we humans manufacture, of many truly causal processes of much shorter duration each.

But there are further confessions to be made as well.

4. ONLY ONE CAUSAL PROCESS TO A BIFURCATION

In a reaction to Papineau's proposal to do away with negative causation, Phil Dowe (1993) presented an example drawn from real life cases in nuclear physics, as a means of revealing what he took to be an undesirable consequence of the proposal. I wish to present a generalization of his example, which will serve to illuminate one source of strain in the relationship between probabilists and causal processes. I shall be presenting example *structures*, that admit of being interpreted in a variety of ways. I shall be interpreting them here, for the sake of definiteness, as stochastic decay events, and interpreting the probabilities (of course) as objective chances of these decay events occurring. But nothing much rests with this interpretation, so long as we understand the probabilities displayed as representing objective chances; the reader is invited to substitute the type of interpretation that best engages her sensibilities. The structures I shall be exhibiting are structures of potential events; thus the structure, as such, cannot be observed in any given episode, since it is not composed purely of actual events. In general, only one potential "sequence" can be actualized in any given instance. We construct these structures after observing populations of relevantly similar sequences of events.

Suppose that a certain substance can decay in any of N equally probable ways, and that the result of each decay process is different from the result of every other. Suppose also that there is an intermediate substance produced in each of these decay processes, and that these intermediate substances too are as different from each other, as they are different from the original substance and from each terminal substance. Figure 1(a) depicts this situation. Are any of these decay processes causal? If we accept the proposal that causes raise the chances of their effects, we will say that every single one of them is. For on the proposal that an intermediate event in a causal process must raise the chances of successive events in that chain, a process A-B-C will qualify as causal only if $P(C|B) > P(C|\neg B)$. In the example of Figure 1(a), the process A-E_i-R_i, for any i, will be causal if we accept the proposal, only if $P(R_i|E_i) > P(R_i|\neg E_i)$. Since, according to the example, $P(R_i|E_i) = 1$ and $P(R_i|\neg E_i) = 0$, the necessary condition is satisfied, and so each branch passes the qualifying test for being causal, which the proposal to reject negative causes sets out.

Suppose now that we have exactly the same situation as before – N equally probable processes – but with the single exception that all the end products are the same. This situation is depicted in Figure 1(b). Are any of these decay processes causal on our proposal to reject negative causes?

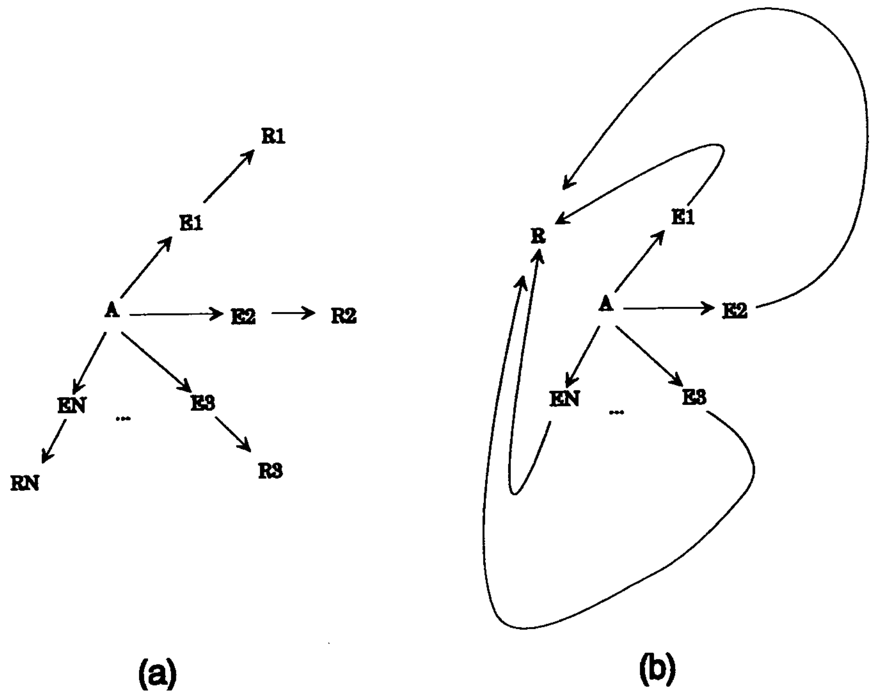

Figure 1.

No, none are, since $P(R_i|E_i) \not> P(R_i|\neg E_i)$, for in that case both chances are equal to unity.

Actually, this is a bit quick. For the proposal to deny negative causes contains no prohibitions against admitting into causal processes events which neither raise nor lower objective probability. Thus if we relax the requirement that events in a causal process must actually raise the probability of the effect, by allowing non-lowering events, we may well be able to accommodate the judgment that the sequences of Figure 1(b) are causal processes too, if those in Figure 1(a) are. But the problem is that this is not just a harmless, friendly amendment to the guiding principle that causes function by raising the probabilities of their effects. In fact, it does *more* harm to our guiding principle, to accept noncontributing events than to accept negatively contributing ones. For we would be admitting events that are, in probabilistic terms, *irrelevant* to production of the effects. (SR, remember, stands for Statistical Relevance.) The probabilist therefore must resist "irrelevant causes" more fervently, as oxymorons, than negative ones. The credentials of a nonlowering event are even more suspect than the credentials of a probability-lowering one, if what counts as making a causal difference, is making a probabilistic difference. It is clear, therefore, that the guiding principle of SR theory is not purely a "formal"

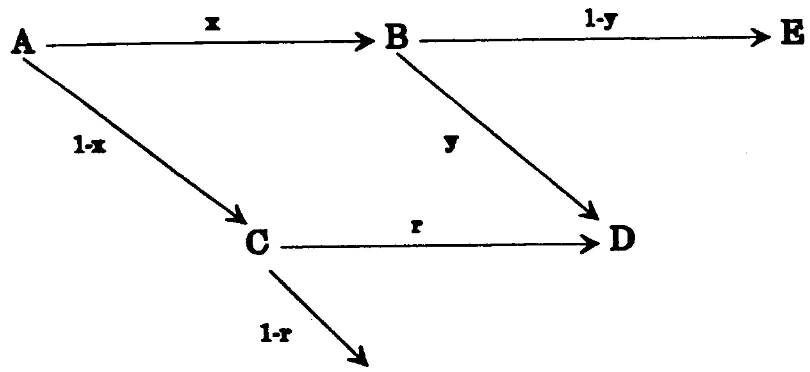

Figure 2.

one; for it were, the nonlowering event would have a higher claim to being included in processes than a lowering event would. And this – the fact that the guiding principle is not purely formal – causes all the trouble, as we'll see. But it suggest that we cannot rely on a purely mechanical rule, such as that all events in a causal process must be chance-raisers for successors, to identify all the events within a causal process.

Suppose now that we have the situation depicted in Figure 2. The process A-B-D is causal, on the hypothesis that causes raise the chances of their effects, only if $y > r$. And A-C-D is causal only if $r > y$. Hence only one path, by which D is brought about from A, can pass the qualifying test. Dowe's example is a special case of the structure of Figure 2, and he argues that it is unacceptable to hold that one process is causal and the other not, particularly if A-B-D represents an alpha decay followed by a beta decay, and A-C-D represents a beta decay followed by an alpha decay. The mere order of the two processes' appearance should not make a difference, or so he thinks, to whether the processes in which they enter qualify as causal.

The avowed probabilist may (unlike Dowe) find this result acceptable, if for example he conceives of causation primarily as a way of coding for the concept of control, or effective strategy, and embraces the idea that an effective strategy is something like the BEST means to a desired end, of which there can of course be at most one.[6] What seems somewhat peculiar here is that *neither* process will qualify as causal if r is exactly equal to y, but one process exactly will qualify if there is the slightest difference between them. What is troublesome is that there is such a very fine line between qualifying and not qualifying. But given that the probabilist's notion of causality is already quantitative, this is scarcely surprising. A purely quantitative notion will surely be insensitive to the distinction between the purely random and the purely stochastic, if that distinction can survive scrutiny.

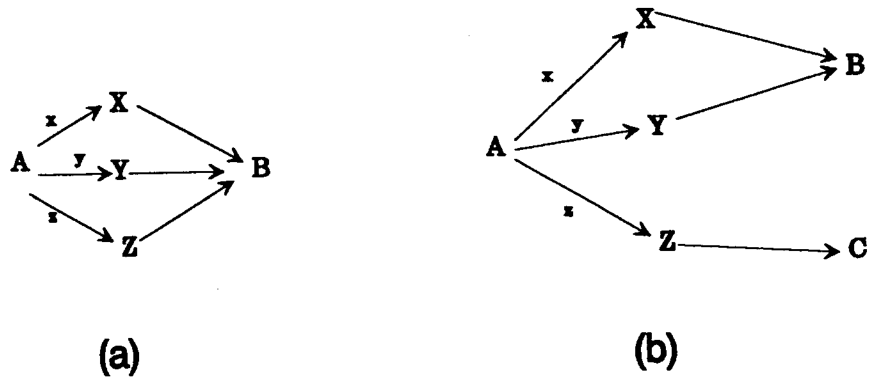

Figure 3.

5. THE MORE DETERMINISTIC THE WORLD, THE LESS CAUSAL ITS PROCESSES

What is most intriguing in this connection, about such structures as I've been displaying, is the pair of structures depicted in Figure 3. Figure 3(a) is an instance of the structure depicted in Figure 1(b), but in which there are just three branches. On the proposal that causes raise the chances of their effects, none of the three processes there represented are causal, no matter what the magnitudes of the objective chances x, y and z. Now Figure 3(b) differs from Figure 3(a) only in that one of the branches from A leads to a different conclusion of the affair than any of the other branches. But due to this difference, *all three* processes of Figure 3(b) pass the qualifying test. This is because the probability of B occurring if A does, in Figure 3(b), is no longer unity. Hence, for example,

$$P(B \mid X) > P(B \mid \neg X) \text{ if and only if } 1 > \frac{y}{y+z}.$$

The inequality on the left hand side of the biconditional is true so long as $z > 0$. Hence all the processes depicted in Figure 3(b) pass the qualifying test so long as all of x, y and z are nonzero.

What's interesting here is this: Figure 3(b) differs from Figure 3(a) in that there is less determinism in Figure 3(b) than in Figure 3(a). We can see this as follows. Let us say that there is more *branching* in Figure 1(a) than in Figure 3(a) if $N > 3$. Let us also say that there is more *pinching off* in Figure 1(b) than in Figure 1(a), because more branches converge in Figure 1(b). Now I take it that the total amount of determinism in a world is a function of (at least) two things: (1) the amount of branching in that world; and (2) the amount of pinching off. A perfectly deterministic world manifests no branching at all. However, pinched off branching contributes less than unpinched-off branching, to the amount of indeterminism there is in a world, when the amount of branching in the two cases is the same. This is true because, other factors being equal, pinched-off branching results in lesser "footage" in the tree of possibilities, and therefore probabilities less thinly smeared across potential events.

The structures of Figure 3 then demonstrate that there is a trade-off between the amount of determinism in a world and the number of causal processes (at least, the number of causal processes with extended duration), on the proposal that only those branches are causal whose intermediate events all pass the qualifying test of being causes of their successors. To be sure, the proposal to reject negative causation will imply – just as intended – that a world in which A and B are connected as in Figure 3(a) will be less causal than a world in which the connection between A and B is unmediated as follows: A →B. But it will also imply that a world in which the connection between A and B is as depicted in Figure 3(b), is a *more* causal world than that in which the connection between A and B is as depicted in Figure 3(a), although P(B|A) in the world of Figure 3(a) is unity. Thus it would appear that when we achieve determinism from A to B, by the kind of pinching off that occurs in Figure 3(a), we sustain a loss in the number of causal processes: no causal processes pass from A to B.

A note of caution: we have *not* demonstrated that those who embrace the principle that causes raise the chances of their effects, cannot embrace the idea that effects are determined by *causes*. For even on that principle, A of Figure 1(a) can qualify as causing B, if A's occurrence raises the chances of B from something less than one. What we've shown is that it's one thing for A to cause B, and something different for A to form part of a causal process which culminates in B. For, as we've seen, the quasi-deterministic world of Figure 3(a), in which $A \rightarrow (X \vee Y \vee Z) \rightarrow B$ holds true, is nonetheless a world in which not one of the processes beginning with A and going all the way through to B passes the usual probabilist's qualifying test for being causal, while all of the processes of Figure 3(b) pass the qualifying test, although none of the outcomes occur with probability one. What we have shown through the examples of Figure 3, is that some chains whose culminations occur with probability one, conditional upon some other events in that chain, do not qualify as causal processes, even if the necessitating events qualify as causes of the end products of these chains. Thus it would appear that if we accept the principle that causes operate by raising the chances of their effects, then we shall have to confess that one event can make a second certain, and even qualify as causing it, whilst nonetheless failing to figure in a causal process that produces it. And this is most disturbing, because it places a wedge between two-term causal relations, on the one hand, and causal processes on the other. And all this *precisely* because we seek a probabilistic account of causality, not despite it.

Someone who believes that insofar as an effect is not determined by its antecedents it isn't caused by them either, will no doubt claim that there is exactly as much by way of causal processes present in the world of Figure 3(a) as there is in the world of Figure 3(b), since only the (short) processes X-B, Y-B and Z-B are causal in either world. However this response papers over the fact that in Figure 3(a) B is also determined *by A,* but that none of the processes leading from A to B qualify as causal (on the proposal that causes raise the chances of their effects). The point is simply that there is a logical or conceptual gap between making an event certain or necessary, on the one hand, and raising its chances, on the other, however compatible the two might be in practice. (In Figure 3(a), A makes B certain, so nothing that does not contribute to A can further enhance the chances of B.) And if we embrace the idea that causing consists in enhancing objective chances, we will not always be able to accommodate the idea that making something certain also amounts to entering into a causal process with that something, without adding codicils to our account.

6. A SOURCE OF THE PROBLEMS

Reductionists aspire to give account of all facts of causation in terms of facts about objective chances, and of course they aspire to do the same vis-à-vis causal processes. It turns out there are additional difficulties, as we've just rehearsed, with the latter. Friends of negative causes might simply say that the probabilists, who view the facts of probability as the fundamental ones, and the facts of causality to be entirely dependent on them, are in this bind purely because they've elected an untenable position about what qualifies as a cause. That, in other words, we should view these troubles as a *reductio* of the principle that causes raise the chances of their effects. Thus that probabilities cannot, all by themselves, be a guide to the facts of causality; so that causal facts must be more fundamental. But the probabilists *could* simply retort that it is the friends of causal processes who have caused the trouble, precisely for directing attention to causal processes. For if we just declare that causal processes are artificial, and hence that all the facts of causation are before us when we have before us the world of two-term causal relations, we escape the problems raised above. I will counsel against this retort: I believe reductionists can do better than they have at accommodating causal processes.

Reductionists gravitate towards the strategy of handling causal processes via the idea that causal processes are just streams of events in which every event that belongs to the stream also qualifies as a cause of the adjoining event, causally downstream.[7] But the proposal that causal processes consist of streams of events in which every member qualifies as a cause of another downstream, goes beyond the guiding principle that facts of causation are nothing more than configurations of facts about objective chances. It is more specific vis-à-vis how to reduce certain salient streams of events to structures of objective chance. Papineau, for example, believes that the notion of causal process can de defined in terms of causal chains – as "bundles of causal chains". "Causal chains are a matter of certain earlier properties making certain later properties more likely than they would otherwise be" (Papineau 1989, 342). He thus defines causal processes as ancestrals of the causal relations. And this precipitates the difficulties we have noticed. It is this specific proposal about processes, and not the fundamental reductionist idea that causal facts are facts about objective chances, that explains the problems reductionists are having vis-à-vis causal processes, as I will now explain. The mistake is in the reductionists' strategy vis-à-vis *processes*, not in the approach to reducing causal relations, as such.

The most exemplary example of a causal process, and the example which Bertrand Russell used to enter into the subject, is the persistence of something – which is often represented as its world line in spacetime. He wrote:

> A causal line may always be regarded as the persistence of something – a person, a table, a photon, or what not. Throughout a given causal line, there may be constancy of quality, constancy of structure, or gradual change in either, but not sudden change of any considerable magnitude. (Russell 1948, 459)

So no matter what our view towards the relationship between causes and probabilities, we should be committed to the idea that, if anything counts as a causal process, the world line of a stationary object does. And most emphatically it should do so in a deterministic world. But then we are faced precisely with the problems I raised above, and in a very pure way. For it might well be that intermediate events in the world line of a stationary object on which no net forces are acting, are in overwhelming numbers non-probability raising. So, for example, suppose that the probability of this chair's being here five minutes from now (E), in LaPlace's world, given that it is here now (C), is (of course) 1. And of course that the probability of E, given that the chair is here two minutes from nows (D) is unchanged – still 1. But surely the event of this chair's here now (C) enters into a causal process with E, if anything does, even though a non-probability raising event intervenes. And surely the intervener D is relevant to E, in a way that events transpiring in mainland China throughout this time are not, despite their also being non-probability lowering.

My suggestion to the reductionists, therefore, is to relax the qualifications for something's entering into a causal process, as follows. The event of this chair's being here two minutes from now (D) is non-probability raising vis-à-vis its being here five minutes from now (E). But it has the very distinguished property of *screening-off* the chair's being here five minutes from now (E), from its being here now (C). And D has this property without qualifying as a cause in the strict sense of raising the relevant probability, if we assume that a deterministic relationship obtains between this chair's successive locations in the next five minutes. And D would qualify for the very same reasons even in a nondeterministic world, where the probabilities of all the events in question were different from 1. But D is relevant (in the colloquial sense) to E, in a way that events transpiring in mainland China are not, since *they* are not screeners-off vis-à-vis C and E. I suggest that we allow screeners-off too to qualify as members of causal processes, in their own right, even if they do not qualify as causes. We will find that the predominant events in the life of a causal process fall into the category of non- probability-lowering, non-probability-raising screeners-off. But once we view screening-off as a credential in its own right for an event's entering into a causal process, we might well wish to allow "negative causes" too to qualify, provided they are also screeners-off. In that way we could solve the ball-and-squirrel problem, and answer Dowe's criticism as well. (As follows: In Figure 2, for example, C screens off D from A, since $P(D|A\&C) = P(D|C) = r$. And likewise B screens D from A. So, on the proposal that screeners off too can qualify as members of causal streams, both A-C-D and A-B-D can qualify as causal.) But what we do here will depend on our taste vis-à-vis negative causes, as such. The point, simply, is that reductionists can accommodate all the intuitions they wish to accommodate about causal processes, without doing violence to the guiding principles of reduction, simply by repealing the idea that all events in a causal process must, without exception, raise the chances of events downstream, and appealing instead to the idea that it is enough if they screen off.

If my suggestion is accepted, processes can be defined exclusively in terms of probability enhancers and screeners-off. Such a reductionist definition, which puts screening off prominently at the center of its understanding of causation, can call on the Theory of Relativity for philosophical support, as I shall now explain.

7. Reichenbach and (quantitative) common cause

Hans Reichenbach thought that the Theory of Relativity could be understood largely in terms of the notion of screening off.[8] Screening off is at the center too of that principle, now known as the *Common Cause Principle*, for which he is most famous. With this principle Reichenbach thought he could give a philosophical account of the physical principles behind the Theory of Relativity. In this section and the next, I will argue that he was right. My argument will require that we scrutinize the notion of screening off qualitatively, as well as quantitatively. This will be explained in the next section. In the meantime I will explain the relationship between Reichenbach's principle and the reductionist strategy I outlined in Section 2.

Very roughly, Reichenbach's Common Cause principle states that uniformities and constant conjunctions amongst events cannot be put down to randomness. Somewhat less roughly (but still not completely, since Reichenbach adds further conditions for there being a common cause), it states that if an improbable coincidence of events has occurred, there must exist an event or combination of events – their common cause – which explains this coincidence by serving as cause of each separately and independently. Thus Reichenbach's is a less refined version of – as much as it is a precursor to – the reductionist strategy I outlined in Section 2 above. It is less refined because it does not go on to explicate the relationship of a common cause to its effects, nor recognize that the (local) probabilistic relation of a common cause to its effects may be indistinguishable from the probabilistic relation of a causal intermediary between a cause and another of its effects.

According to Reichenbach's Common Cause principle, it is to be sure *individual* events, and not classes or populations of events, which stand in two-term causal relations to one another. Even so, the Common Cause principle (like the reductionist approach) rests on the philosophically significant observation that it is a *conjunction* or *correlation* amongst events[9] – and not simply some observable relation between, for example, two individual events in sequence – which calls attention to itself, raising thereby a demand for causal explanation. And thus the Common Cause principle, like the reductionist doctrine, gives expression to the idea that it is causes that elevate the probabilities of things occurring. Thus advocates of the Common Cause principle too can view themselves as in agreement with Hume when he propounds, as a good empiricist, the proposition that causal relations can be neither observed nor grasped by intuition, but must be inferred, and only after considerable experience with the types of events in question. And that, furthermore, the inference from experience to cause is thoroughly nontrivial. They might disagree with Hume's pessimism, however, and proclaim that the inference in question can legitimately be made, with some considerable hard work.

So far, then, the Common Cause has everything in common with the reductionist approach. Here, now, is why we should, with Reichenbach, view the reductionist enterprise as boosted by the Theory of Relativity.

8. QUALITATIVE COMMON CAUSE: THE "MEANING" OF SCREENING OFF

The now-familiar treatment of causality via probabilistic analysis, whether it aspires to reduction or not, is nonetheless quantitative: it is conducted purely in terms of numerical magnitudes, and in terms also of their equalities or inequalities. This quantitative analysis fails to illuminate the qualitative features of the relation of causal dependence being articulated through the notion of screening off. The quantitative treatment is concerned with the details of specific two-term relations, rather than with the universal characteristics of the two-term relation wherever it is manifested. It fails to illuminate the structural and global features so analyzed. These qualitative, universal features, once grasped, will enable a deeper appreciation of the relation between probabilistic causality, on the one hand, and the Theory of Relativity, on the other. To this end I shall call on the very elegant axiomatic treatment of common cause reasoning formulated by Judea Pearl in his very comprehensive work *Probabilistic Reasoning in Intelligent Systems*.[10] I will cast Pearl's axioms in terminology that is intended to make the transition to Relativity most transparent. For we shall require language capable of handling common causes and effects that are sets of quantities with continuous or discrete sets of values, in addition to cases in which we have simply nameable single events that occur or do not occur.

A *quantity*, I shall say, is a characteristic of the universe which may change in magnitude over time. (A *property*, by contrast, is a definite magnitude of a quantity: when the position of a body B is x, then B exhibits the property of *being located at x*.) An *event*, on the other hand, is the occurrence consisting of a quantity taking on a definite magnitude (at a definite time). A *process,* consequently, will be a collection of quantities taking on magnitudes together, such that the alterations to these quantities occur in a continuous fashion. (This last definition accords with Russell's notion of process, discussed above in Section 6.)

We wish now to use the notions of *quantity* and *event* to give precise expression to the notion of causal *dependence*, á la Pearl, in reductionist terms. Let X, Y, and Z represent nonoverlapping sets of quantities: $X_1, X_2 \ldots, Y_1, Y_2 \ldots$, and $Z_1, Z_2 \ldots$; and let the symbol $\cup$ represent unions of quantity sets. Then let E_Q^q represent the event or events consisting of quantities in the set Q taking on *magnitude configuration* q, which consists of magnitudes drawn from the domains of corresponding quantities, one from each. So $E_{X_1}^{x_1}$, for example, will represent the event that quantity X_1 takes on magnitude x_1 in the domain of X_1; while E_X^x, or simply E_X, will represent the complex of events of X_1 taking on magnitude x_1, and X_2 taking on magnitude $x_2, \ldots$.

Now let $CI(E_Q^q, E_S^s, E_R^r)$, or simply $CI(E_Q, E_S, E_R)$, stand for E_Q^q *is causally independent of* E_R^r*, conditional upon* E_S^s. Essentially, this will correspond to the notion of screener-off, introduced as Definition 2.1 above, and hence will hold exactly on condition that E_Q^q, E_R^r, and E_s^s satisfy:

DEFINITION 8.1 (Screening off). $CI(E_Q^q, E_S^s, E_R^r) \equiv_{df} P(E_Q^q | E_s^s \wedge E_R^r) = P(E_Q^q | E_s^s)$.

So for example, $CI(E_X, E_Z, E_{Y \cup W})$ will hold exactly on condition that $P(E_X | E_Y \wedge E_W \wedge E_Z) = P(E_X | E_Z)$. Since $CI(E_X, E_Z, E_Y)$ holds exactly when E_X and E_Y are uncorrelated once contributions, to each, from E_Z have been accounted for, *un*conditional independence will be denoted by $CI(E_X, \oslash, E_Y)$. And causal *dependence* will be the absence of causal independence, and will have to be diagnosed indirectly – via some strategy such as is embodied in A1–A3.

Pearl (1988) shows that CI must satisfy the following four logically independent conditions:

- *Symmetry*: $CI(E_X, E_Z, E_Y) \Longleftrightarrow CI(E_Y, E_Z, E_X)$
- *Decomposition*: $CI(E_X, E_Z, E_{Y \cup W}) \Longrightarrow CI(E_X, E_Z, E_Y) \& CI(E_X, E_Z, E_W)$
- *Weak Union*: $CI(E_X, E_Z, E_{Y \cup W}) \Longrightarrow CI(E_X, E_{Z \cup W}, E_Y)$
- *Contraction*: $CI(E_X, E_Z, E_Y) \& CI(E_X, E_{Z \cup Y}, E_W) \Longrightarrow CI(E_X, E_Z, E_{Y \cup W})$

And if no magnitude of a quantity is impossible (if, that is, no magnitude of a quantity receives a probability assignment of zero), then a fifth condition must hold as well:

- *Intersection*: $CI(E_X, E_{Z \cup W}, E_Y) \& CI(E_X, E_{Z \cup Y}, E_W) \Longrightarrow CI(E_X, E_Z, E_{Y \cup W})$.

These axioms present the notion of screening off qualitatively, universally and globally. They turn our attentions away from the arbitrary choices that are made when quantification of the relation of causal dependence is made – which, to be sure, is absolutely necessary and perfectly legitimate when the aim is to compare different strengths of causal dependence. However the arbitrary choices, together with application of the axioms of probability, put the philosophical substance of the idea of screening off, completely out of focus.

Now Pearl and Paz (1985) conjecture that the first four axioms are complete when 'CI' is defined as above – as follows: for each three-term relation satisfying the four axioms, there exists a probability model P such that $P(E_x|E_Y \wedge E_Z) = P(E_X|E_Z)$ *iff* $CI(E_X, E_Z, E_Y)$. I will simply accept this completeness conjecture as true, on the palpable authority of its authors. (A thorough treatment of the completeness problem and related results is found in Geiger and Pearl (1988).) What, now, is the philosophical payoff?

Pearl (1988, 85 ff.) demonstrates that the axioms governing CI can also be construed as articulating the notion of spatial separation, or disconnection. In other words, that the concept of independence which (Screening off) expresses, on the one hand, and the concept of spatial disconnection on the other, are – strikingly – qualitatively indistinguishable. For the axioms may be depicted as in Figure 4, and construed as follows: *Symmetry* states that if Z separates X from Y, it also separates Y from X. *Decomposition* states that if Z separates X from a composite Y, then it separates X from every subcomponent of Y. *Weak union* sets out a condition in which a separating region can be nontrivially compounded and still continue to separate; the condition is that the augmentation should come from the region initially separated by that which will continue to separate. *Contraction* states that a separating region may be reduced, if the deleted portion is also separated from one of the regions initially separated. And *Intersection* states that if, among the subregions of a certain region, X can be separated from the rest by two different subregions of the whole, then the intersection of the two different subregions will also separate X from the rest of the whole.

Thus, if we accept Pearl's completeness conjecture, we will be in a position to view (Screening off) as stating very simply that spatiotemporal separation is a barrier to causal interaction. For accepting the completeness conjecture just amounts to saying that there is no more to the configuration of relations CI than a structural, qualitative reality that is captured by the axioms of spatial separation.

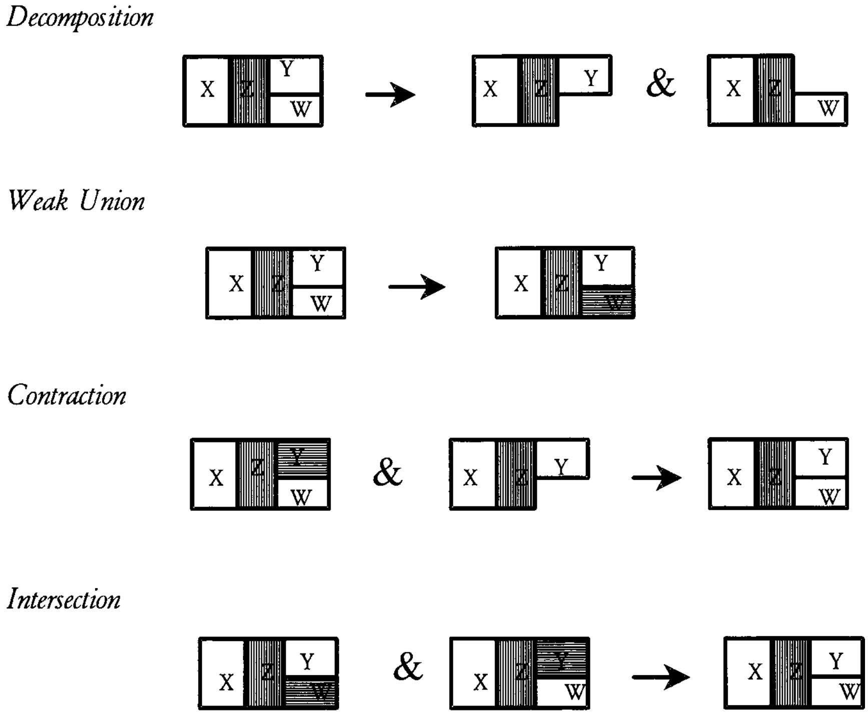

Figure 4.

True: (Screening off) says nothing about there being an asymmetry between cause and effect, whereas one would expect an account of causal dependence to do so. And the reason, quite simply, is this: that (Screening off) is concerned, not with causal DEpendence, but with causal INdependence. For it is intended as only the focal center of an approach to diagnosing causal dependence, indirectly, through assessments of causal independence. Thus, (Screening off) requires being surrounded by statements of that strategy, such as for example we have in A1–A3. And of course causal independence is symmetrical, whereas causal dependence is not. So whereas I think it is absolutely correct to say that the principle (Screening off) lies at the heart of every empirical strategy for discriminating causal relations, I think it is also absolutely clear that it cannot be all there is to the *concept* of causation. Essentially, common-causal reasoning embodies an indirect *approach* to causal analysis: it eliminates what cannot possibly be directly causally dependent, and leaves it up to us to complete the picture in (at least so we hope) the only way we can – a procedure Sherlock Holmes might have commended. And so it should come as no surprise at all that (Screening off) implies no asymmetries, or that its axiomatic representation is no different from that of (spatial or spatiotemporal) separation. And critics of the reduction enterprise are right in stating that, if all there is to the probabilistic analysis of causality is (Screening off), then probabilistic causality is a poor substitute for the real thing. For every reductionist accepts that (Screening off) can't be the whole story, even if they embrace it as the story's leading figure. The reductionist story must include the means of inferring causal dependences from constellations of causal independences. Reductionists will thus be required to supplement (Screening off) with a story about how to judge, of two events that turn out to be neither conditionally nor unconditionally independent, how the one stands, whether as cause or as effect, to the other. And this is precisely what they have undertaken to do, as I described in Section 2.

The problem for probabilists, then, is to show, or at least to make it plausible, that what the indirect procedure leaves undone, can only be done one way. That is, that there is only one way to complete the picture that is left incomplete once all the independence data that can be collected has been collected. It is a monumental position to defend. And all the complaints of the critics are, in one way or another, variations on the theme that it can't be done. This is where the difficult work of Spirtes, Glymour and Scheines (1992) comes in.

What, now, follows from our new-found, qualitative understanding of (Screening off)? Just what Reichenbach suspected: that (Screening off) captures everything behind the famous principle of no action at a distance, in which Einstein believed firmly, and which played a central role in his formulation of relativity. Here is why.

9. COMMON CAUSE VIS-À-VIS ACTION AT A DISTANCE

Einstein, famously, held that there can be no *action* (using the term interchangeably with 'cause') at a distance, and built the theory of relativity up around this principle. What was his conception of action, and what his rationale for the principle that it could not take place between separated things? Einstein appears to have founded his conviction in the principle of no action at a distance upon two other principles, which themselves are prompted by the results of experimental observation: (1) that the speed of light in a vacuum is fixed and independent of the velocity of its source in relation to any reference frame from which the light signal can be viewed;[11] and (2) that there is no absolute *same-time* relations among events – that, instead, the only absolute relations involve "spacetime" intervals, rather than time intervals purely.[12] The second principle, which normally is referred to as the *relativity of simultaneity*, rests – as Reichenbach brilliantly argued – on the nonabsolute character of velocity itself. In other words, the relativity of simultaneity is a consequence of the fact that measurement of velocity presupposes what Reichenbach called a "coordinative definition" of simultaneity (Reichenbach 1958, 123–129).

The two principles on which Einstein called are not unrelated: for it is the principle of relativity of simultaneity which allows us to take the speed of light as fixed rather than as varying, without falling into contradiction with any potentially observable phenomena. And rather than these two principles *entailing* Einstein's proposition that there can be no causal action (understood very broadly as a signal or a transfer of energy or information) at a distance, these principles instead *presuppose* it. For they presuppose, as Reichenbach was quite right to point out, the principle according to which the world line of an object is *ordered* as to causal influence. Thus the principle that no action can occur at a distance is quite metaphysical in character, as we might say. What does it involve?

The principle that action does not occur at a distance answers two important questions: (1) what are the *units* in which happenings in the world break down, if we care to give a causal account of relations between them? and (2) what sort of thing is causal influence, and how does it flow from the causing event to its effects? The answer to the first question is what may be referred to as a *spacetime* answer: that for purposes of causal analysis events are localizable at unextended spacetime points.[13] And the answer to the second question, according to the principle of no action at a distance, is quite simply that causal influence is transmitted or transferred via contact, as a parcel may be passed about: the flow of causal influence involves a continuous sequence of asymmetrical transactions among events. Hence causal influence is an all-but-palpable thing, according to Einstein's conception, which passes along spacetime trajectories much as bodies themselves pass. One is even tempted to say of this conception that it comprehends causal influence as paradigmatically *exemplified* by bodies, and hence that it is a mechanistic conception of causation. Officially speaking, then, there can exist no finite spatiotemporal distance between a causing event and that which it brings about directly (i.e., without going through intermediaries); but we may relax this restriction when it comes to causes and their *mediated* effects. Thus on this conception, if some event A has had causal impact on some event B separated from A by an expanse of spacetime, then causal influence must have traveled across that expanse. It follows that when no influence passes from A to B through the intervening expanse of spacetime, possibly because the upper bound on signal transmission prohibits it, then there can be no causal dependence between the two events.

The fact that Pearl's axioms – which I am taking for granted as extracting all the philosophical content of (Screening off) – can be interpreted as enunciating the principles of spatial separation or disconnection is both suggestive and penetrating, particularly if we stipulate that the entities which enter causal relations – events – are ones that inhabit a spacetime structure at points, and that this structure is representable geometrically. Pearl's insight therefore grounds a quite concrete argument for the proposition, taken for granted by Einstein, to the effect that spatiotemporal separation is *prima facie* grounds for causal independence. If we identify events with their spacetime locations, then this becomes just another way of saying that there can be no causal action at a *spatiotemporal* separation. Spatiotemporal separation, in other words, is a barrier to causal interaction – but one which is surmounted by movement of causal influence across the intervening separation. This notion of spacetime separation as a barrier to interaction is precisely the notion which forms the basis of relativity theory. Relativity treats influences as carried along "world lines" or trajectories in spacetime, and treats "space-like" separated events differently from "time-like" separated ones. Space-like related events – defined as events taking place at separated spacetime locations in such a way that light cannot travel from the one's position to the other's in the time intervening between the two events – are universally taken as causally independent; whereas time-like separated events – defined as not space-like separated – are in the market for causal dependence, and become causally connected when a so-called "causal signal," traveling no faster than light, passes from one to the other.

It pays to be somewhat more pedantic about this argument, in support of what is entitled to the name of a *transference conception of causation*. Suppose we begin, like Einstein, with the assumption that spacetime separation is *prima facie* grounds for causal independence. As I'm urging, on the authority of Pearl's axioms, this is no more and no less than the philosophical kernel behind the principle (Screening off). When presented with two spatiotemporally separated events, we first diagnose the possibility of causal dependence by inquiring whether there is a relevant correlation. If there is not, a search for causal explanation is unwarranted. (This does not mean that the events in question are causally *irrelevant* to each other, as both negative and positive causal contributions from one event to another may cancel out to zero.[14] Causal *dependence*, however, is the net sum of negative and positive contributions. So there may be causal relevance without causal dependence.)

If, on the other hand, there is a relevant correlation, this (according to the reasoning in A1–A3) requires causal explanation, in the form of some third event against which the original two are *conditionally* independent. On the orthodox view, this third event must lie either in the spatiotemporal expanse "between" the original two, or else in their common past.[15] On a counterorthodox view, but still one which embraces the rationale behind common- causal reasoning, the third event need not lie "backwards" of either of the original events, as causal influence can flow either forwards or backwards in time. In either case, causally dependent events can be space-like related only if they are conditionally independent with respect to an event to which each is time-like related, since, by hypothesis, spatiotemporal separation is *prima facie* grounds for causal independence.

This process of revealing intermediaries (screeners-off) in relation to which an original statistical dependence is reduced to a conditional independence, leads by the axioms, to construction of continuous chains of causal connections, along which, as we we might be inclined to say, causal influence *travels*. These are the processes, if you like the term. Thus if we embrace (Screening off) plus A1–A3, we shall be able to say that causal dependence between separated events, whether the separation is time-like or space-like, requires a movement of influence across expanses of spacetime. And that these are the causal processes. In that way we can demonstrate that probabilistic insights, in the form of (Screening off), lie at the heart of a spacetime approach to causal processes.

Thus reductionists can claim title to their idea of (Screening off) as the very same as that idea, pioneered by Isaac Newton and improved upon by Einstein, that change in the physical *status quo* occurs only when an *influence* (in the form of matter or energy or momentum or charge or something "active" like these) brings it about; that influence is in motion, very much as a body itself can be in motion; and that influence *flows* from one event to another – not unlike the way a parcel might be passed from one hand to another – in a continuous series of transactions between spatiotemporally contiguous pairs of (spatiotemporally ordered) events. By standing behind Pearl's conjecture, reductionists can claim that their conception of causal independence shares its qualitative character with that of spatial separation. And thereby reductionists can, through the argument just rehearsed, lay claim to giving an adequate treatment of processes. The arguments casts doubt on views (for example, Irzik (1996); Salmon (1984); and Sober (1988)) according to which the conception of causal process is in some way more fundamental than that of objective chance and therefore ineliminable, by reduction, from any respectable account of causation.

10. TWO POINTS OF POTENTIAL TROUBLE

There are two points of potential trouble for my proposal in section 6, via screeners-off, for handling the problems facing probabilists trying to reduce causal processes. The first is that screeners-off are not always causal intermediaries; sometimes screeners-off are joint effects of a common cause. So for example, if the chair's being here now (C) is a common cause both of its being here 5 minutes from now (E), and also of certain light beams racing off into space two minutes from now (D), and if, furthermore those light beams could not have got where they did except by the chair's being here now just as it is, then D screens off C from E. But surely we don't wish to say that it participates in the causal process consisting of this chair's presence here for the duration of 5 minutes. Thus being a screener-off is not a sufficient credential for belonging in a causal process together with the incidents which it screens off.

Agreed. But there is a very simply reply for probabilists here. The point I'm trying to make is simply that the events which belong in a causal process, but which are not probability-enhancing, are causal intermediaries set *between* probability-enhancing ones. (Of course this idea, which emerges from thinking about how events in a deterministic world might still qualify as causal, does not apply strictly to such a world, since in a deterministic world there are no probability enhancers at all.) Call the events which don't enhance the probabilities of events they connect, but which nonetheless belong in the causal process, non-enhancers. The proposal I made in Section 6 can then be put as follows: *non-enhancers in a causal process are nonetheless causal intermediaries between enhancers.* As I have been emphasizing, it is a nontrivial empirical exercise (and possibly unworkable or even in principle impossible) to determine whether a screener-off is an intermediary or merely a joint effect, with one of the events screened, of a common cause. Now the question whether some event is an intermediary or a joint effect, in relation to some other event, depends not *only* on its precise numerical contribution to that other event, but also on whether it is surrounded in spacetime exclusively by non-enhancers. I said that we should regard non-enhancers in a causal process nonetheless as causal intermediaries set *between* enhancers. Thus whether a non-enhancer falls in some stream of events that deserves calling causal, is a non-local affair. It depends on *more* than its own numerical contribution to production of other events in that stream. Nonetheless, the more on which it depends, is just more of the same sort of thing: more along the lines of probability-enhancing facts. Thus whether a non-enhancer falls in some stream of events that deserves calling causal, is a non-local affair, but it is nonetheless a probabilistic one.

So consider just a simple case. We discover that A enhances the probability of B, but that S, a temporal intermediary, screens B off from A. We eliminate the possibility that S cannot be a common cause, and thus are left with two possibilities: S is either a causal intermediary or an effect (joint with B) of A. Suppose we can determine it is the former. We have not yet answered whether S enters into an unbroken causal process with A and B; and in particular, we have not determined whether S is itself an enhancer of B. On the familiar approach to processes I am criticizing, namely that to belong to a causal process something must qualify as an enhancer, S does not belong. I am suggesting that this consideration should not be decisive. That being intermediary is sufficient as a qualification, whether one also enhances or not. So as I announced at the outset, if probabilists can satisfy us that their methods can resolve the question: Given any two events, what is the causal relation between them? then they can satisfy us too on the subject of causal processes. And I am here simply taking the antecedent for granted.

I am taking for granted, as announced at the beginning, that probabilists can identify each of the intermediaries and joint effects for what they are, to satisfaction, so as to be in a position to draw the directions of the arrows they seek to draw in order to represent the two-term causal relations of the form *C-causes-E*. I am assuming, as announced at the beginning, that probabilists are in a position to say, of any pair of events, whether there is a relation of causation between them (that they are in a position to say that intermediaries and joint effects come out on different sides of the divide because intermediaries *might* be bona-fide causes, but joint effects are *definitely* not). Now if probabilists *cannot* make these discriminations, systematically and unambiguously, at least under the very best of conditions, then they're in trouble all the way down, and not just with accommodating our notions about causal processes.

The problems vis-à-vis processes with which probabilists have to deal, are *not* to the effect that probabilists fail to discriminate correctly how a certain event contributes to the production of a certain other. It is in fact taken for granted that everyone can agree on the positive/negative contribution facts, and yet disagree as to whether the event streams in question are causal. When it comes to negative causes and their cousins, what is in dispute is whether a certain sequence of episodes comprises a single, continuous causal process, or whether a would-be causal process is broken by the presence of a chance-reducing event. So the controversy I have entered on the probabilists' side, is not over whether the probabilist has correctly assessed the precise contributions of certain events to the production of others. It is instead over how to categorize streams of events, dividing between causal and noncausal. And thus a contention to the effect that probabilists will misdiagnose the numerical contribution of a cause to an effect is out of place in this setting.

The second point of potential trouble for my account is just the converse of the first: causal intermediaries are not always screeners off. Consider the case of two radioactive atoms, secluded in a chamber. At t_0 we apply a burst of stimulating radiation; call this event C. In the absence of stimulating radiation, there is still some low probability of either atom decaying, but the chances of decay are enhanced by the burst of stimulating radiation. At t_1 one of the two atoms – atom A – decays; this is event D_A. Then at t_2 our detector clicks; let E name the event that the detector clicks at least once. Let's suppose that the detector is not infallible, and hence that $P(E|D_A)$, like $P(E|D_B)$, is somewhat less than 1. We wish the sequence C-D_A-E to come out as a causal process. But for some rates of detector fallibility, D_A may not screen C off from E. For, if $P(E|D_A) < 1$ by a sufficient amount, $P(E|D_A)$ can be considerably less than $P(E|C\&D_A)$, since C (but not D_A) enhances the probability that the second of the two atoms – atom B – decays. So D_A, causal intermediary between C and E, need not in every instance screen the one off from the other.

I accepted the counterexample for the first point, and refined my account to accommodate it, by drawing attention to non-local facts of enhancement and non-enhancement. But I shall not do it for the second point. I cannot, on behalf of probabilists, accept that causal intermediaries are not screeners off: it is a core principle for inferring from probabilities to causes. Thus I will exhibit a strategy for rejecting this, and all proposed counterexamples like it. The strategy essentially involves the fact that we have taken a spacetime approach to events.[16] My strategy will then be to identify the suitable spacetime regions which causal intermediaries screen off from each other. And I will suggest that it is always possible to do it in the face of such proposed counterexamples as the one we have just rehearsed.

I therefore ask you to consider the proposition that C-D_A-E is not, strictly speaking, a bona fide causal process. Because it includes too much. I suggest that the causal process from D_A to E is not inaugurated by C, as such, but rather by C_A, the event of atom A being excited with a burst of stimulating radiation. Thus the causal process, in which D_A enters as a causal intermediary is C_A-D_A-E, and in that process it is without question a screener off. Thus the strategy I am proposing is to narrow or enlarge the identified spacetime regions, in proposed counterexamples, so that the intermediary will function as a screener off. I suggest this will always work.

You might reply that if my strategy for rejecting the proposed radiation counterexample is acceptable, then we should have saved ourselves the trouble and simply rejected Dowe's initial counterexample earlier on. I reply that one can be moved by Dowe's counterexample but not by this one, particularly if one's orientation towards the analyzing of basic units of causal interaction is – as Einstein's was – in terms of spatiotemporal regions rather than linguistic or logical conglomerations. Thus, if one's concerns are relativistic, one will have quite clear motivations for rejecting the proposed radiation counterexample, which do not apply to Dowe's.

But you might say that in rejecting the proposed radiation counterexample example we have had to introduce extra-probabilistic considerations, which enable us to say that the cause of $D_A - E$ is not C but C_A, and we have thereby compromised our true reductionist aspirations. We have had to settle for less than the original reductionist proposal. This charge, I believe, is mistaken. For the approach to events we have already settled upon – the approach which makes reduction an aspiration in the first place – is an approach that identifies events with spatiotemporal localities. And the strategy reductionist have elected – the strategy which puts screening off at the center – amounts simply to saying (as Pearl's work abundantly illustrates) that the *size* of the spatiotemporal locality matters. Because spatial separation, as such, is *ex ante* evidence for causal independence. Thus attention to the spatio-temporal is not extra-probabilistic. It is what motivates our strategy in the first place.

Thus I conclude that causal process pose neither extra, nor special work for probabilists, over and above what already they have to do to satisfy us on the subject of causal relations. But this, by itself, is not all good news. It is a double-edged point. For it focuses attention once more upon the difficult question of whether the two-term causal relation can be handled purely as a configuration of probabilistic relations. The news, therefore, is that when it comes to the reduction of causation, we should not divert attention from the fundamental question of whether the reduction of the two-term relation 'C causes E' is satisfactory.

ACKNOWLEDGEMENTS

For directing me to the work of Judea Pearl I thank Barry Smith. For insightful remarks and criticisms I thank Christopher Hitchcock and several anonymous referees. And for that which remains dark still, despite the help I have received, I have only myself to thank.

NOTES

[1] Papineau adds the following condition as well: $P(E|C\&\neg S) = P(E|\neg S)$. For our purposes this clause is unnecessary, as we shall be dealing in the end only with sets of actual events.

[2] Probably the most famous example of this sort is due to Good (1961). It is the Holmes-Moriarty example.

[3] This is as a matter of fact incorrect: not everyone will prepared to grant this as a case of negative contribution, but some will make an effort to account of all the events of the process as each one making a successor more likely. One might, as for example Christopher Hitchcock suggested in personal correspondence, follow Deborah Rosen's lead (as reported in Suppes (1970, 41)) in suggesting that the squirrel kick raised the probability of a successor event, and that that successor raises the probability of the ball going into the cup. So suppose the golfer's putt sends the ball rolling, it is kicked by the squirrel in such a way as to encounter a rock, which changes its trajectory so that it falls into the cup. The initial putt increases the probability of the squirrel kicking it; the squirrel's kick raises the chances that it will encounter the rack just as it did; and the encounter with the rock, in just the fashion it did, raises its chances of making its mark. This strategy will not always be successful: it will fail in cases I introduce in the next section – example *structures* (not stories as such) specifying specifically what the probability increments are from one node in the structure of events to the next. They have been shown to have real-life instances, but even if they had not, they would still cause real-life trouble for the probabilist. And that is why these structures were invented.

[4] Salmon (1984); see also the discussions of the mark criterion by Kitcher (1989) and Dowe (1992).

[5] See the transference theories of Aronson (1971) and Fair (1967), criticized by Dowe (1993).

[6] For example, Mellor (1988) might not find it unacceptable.

[7] Of course this might be taken as cheating, because it assumes we can get away with talking about events as adjoining, even though time is not discrete. I think it's not really cheating, because we can think of events, in general, as occupying space-time regions rather than unextended space-time points. But this is not the place to discuss the matter. And anyway Papineau (1989, 331–334) has a work-around, if you prefer a doctrine of events as occupying only space-time points.

[8] In *The Direction of Time* (Reichenbach 1956), Reichenbach sought to explain why in certain time-ordered sequences of events there exists an objective *future orientation* (i.e., *before-and-after* relations) in terms of causal dependence relations amongst these events. Reichenbach, in other words, strove for a reduction of temporal direction to the direction of causation, conceiving of the latter relation as one which obtains between individual events taken in pairs, and consequently capable of forming chains. The task he set for himself was that of articulating a conception of cause, founded on considerations from fundamental physical theories, which would ultimately admit of a directionality.

[9] The term *correlation* comprehends irregular or imperfect conjunctions, as well as conjunctions of the constant variety.

[10] I will not be concerned here with the systematic advantages of the axiomatic treatment of causal dependence and information. These, as well as matters concerning the virtues and advantages enjoyed by systems which adhere to these axioms in organizing information and drawing inferences, are also superbly treated by Pearl. Efficiency is perhaps

responsible for prevalence, since efficiency contributes to advantage over competitors. But of course the question whether causal models are efficient is distinct from the question whether they portray the world accurately. For a causal model may be both efficient and false.

[11] This proposition receives confirmation from observations made of double-star systems.

[12] The discussions at Geroch (1978, 67–91) and Reichenbach (1958, 123–129) are helpful for illuminating these principles and the experimental bases on which they rest.

[13] This was suggested to me by Butterfield (1989), although he does not spell out the matter as I do here.

[14] Certain of Cartwright's (1989) examples reveal as much, but no concrete examples are necessary for the point.

[15] Cf. the discussion at in Papineau (1993, 240 ff.).

[16] See also Butterfield (1989).

REFERENCES

Aronson, J.: 1971, 'On the Grammar of "Cause" ', *Synthese* **22**, 413–430.

Butterfield, J.: 1989, 'A Space-time Approach to the Bell Inequality', in *Philosophical Consequences of Quantum Theory*, Notre Dame University Press, pp. 114–144.

Cartwright, N.: 1989, *Nature's Capacities and their Measurement*, Clarendon Press, Oxford.

Dowe, P.: 1992, 'Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory', *Philosophy of Science* **59**, 195–216.

Dowe, P.: 1993 'On the Reduction of Process Causality to Statistical Relevance', *British Journal for the Philosophy of Science* **44**, 325–327.

Dowe, P.: 1995, 'What's Right and What's Wrong with Transference Theories', *Erkenntnis* **42**, 363–374.

Eells, E. and E.Sober: 1983, 'Probabilistic Causality and the Question of Transitivity', *Philosophy of Science* **50**, 35–57.

Fair,. D.: 1967, 'Causation and the Flow of Energy', *Erkenntnis* **14**, 219–250.

Geiger, D. and J. Pearl: 1988, 'Logical and Algorthmic Properties of Conditional Independence', Technical report, Cognitive Systems Laboratory, UCLA.

Geroch, R.: 1978, *Relativity from A to B*, University of Chicago Press, Chicago.

Good, I. J.: 1961, *A Causal Calculus*, i–ii.

Irzik, G.: 1996, 'Can Causes be Reduced to Correlations', *British Journal for the Philosophy of Science* **47**, 249–270.

Kitcher, P.: 1989, 'Explanatory Unification and the Causal Structure of the World', in P. Kitcher and W. Salmon (eds), *Minnesota Studies in the Philosophy of Science*, University of Minnesota Press, Minneapolis.

Mellor, D. H.: 1988, 'On Raising the Chances of Effects', in J. H. Fetzer (ed.), *Probability and Causality*, Kluwer, Dordrecht, pp. 229–39.

Glymour, C., Spirtes, P. and Scheines, R.: 1992, *Causation, Prediction and Search*, Springer-Verlag.

Papineau, D.: 1989, 'Pure, Mixed and Spurious Probabilities and Their Significance for a Reductionist Theory of Causation', in P. Kitcher and W. Salmon (eds), *Minnesota Studies in the Philosophy of Science*, University of Minnesota Press, Minneapolis, pp. 347–388.

Papineau, D.: 1993, 'Can We Reduce Causal Direction to Probabilities?', *Proceedings of the 1992 Bienniel Meeting of the Philosophy of Science Association* **2**, 238–252.
Pearl, J.: 1988, *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufman, San Mateo.
Pearl, J. and Paz, A.: 1985, 'Graphoids: A Graph-based Logic for Reasoning About Relevance Relations', in D. Hogg B. DuBoulay and L. Steels (eds), *Advances in Artificial Intelligence 2*, North-Holland, Amsterdam.
Reichenbach, H.: 1956, *The Direction of Time*, University of California Press, Los Angeles.
Reichenbach, H.: 1958, *The Philosophy of Space and Time*, Dover, New York.
Russell, B.: 1948 *Human Knowledge*, Simon and Schuster, New York.
Salmon, W. C.: 1984, *Scientific Explanation and the Causal Structure of the World*, Princeton University Press, Princeton.
Sober, E.: 1988, 'The Principle of the Common Cause', in J. H. Fetzer (ed.), *Probability and Causality*, Kluwer, Dordrecht.
Suppes, P.: 1970 *A Probabilistic Theory of Causality*, North-Holland, Amsterdam.

Department of Philosophy
University of Utah
341 Orson Spencer Hall
Salt Lake City, UT 84112-1107
U.S.A.
E-mail: mariam.thalos@philosophy.utah.edu