[forthcoming in a special issue of
Synthese on Probabilities and Propensities edited by James H. Fetzer]

PROPENSITY TRAJECTORIES, PREEMPTION, AND THE IDENTITY OF EVENTS*

Ellery Eells
University of Wisconsin-Madison, Philosophy
May 9, 2001

Abstract: I explore the problem of "probabilistic causal preemption" in the context of a "propensity trajectory" theory of singular probabilistic causation. This involves a particular conception of events and a substantive thesis concerning events so conceived.

Key words: causation, event, token, type, preemption, probability, propensity

The problem of preemption for theories of causation is well known. In it's original and basic form, it is a problem for theories of causation that take a cause to be some sort of a necessary condition for its effects (e.g., an INUS condition as in J. L. Mackie (1974), or an event upon which an effect counterfactually depends as in David Lewis (1973, 1986)). The problem is that an event c may very well be a cause of an event e even though, had c not occurred, a "backup" event c' which also occurred would have caused e instead; in this case, c is not necessary for e, and c "preempted" c' in causing e. A clear and simple example of this is described by David Lewis (2000) and by Ned Hall (2001). Suzy and Billy both want to break a bottle with a rock; they are both experts at this; they both pick up rocks and throw, accurately; Suzy's rock reaches the bottle first, breaking the bottle; by the time Billy's rock arrives the bottle is already shattered and his rock hits no glass at all. In this case it is clear that it was Suzy's throw (c) that caused the bottle to break (e), but Billy's throw (c') would have caused the bottle to break had Suzy not thrown her rock. In this case, Suzy's throw was not, under the circumstances, necessary for the bottle to break; her throw preempted Billy's throw in the breaking of the bottle.[1]

---

*I thank Dan Hausman, Christopher Hitchcock, Stephen Leeds, Alfred Mele, and Elliott Sober for useful comments or discussion.

[1]I note here, for reference in a later note, that preemption in a case like this is sometimes called "late preemption", in contrast to cases of "early preemption" in which a causal chain from the

Recently, the problem of preemption has been applied to *probabilistic* theories of causation, on which a cause need not be necessary for its effects, but only raise the probability of its effects. Peter Menzies (1996) has applied the problem to counterfactual theories of probabilistic causation, and Douglas Ehring (1994, 1997) has applied the problem to a "probability (or propensity) trajectory" approach to probabilistic causation that I have earlier urged (1991). In this paper, I will be concerned with applying this trajectory approach to cope with the problem of (what I will call) probabilistic preemption. I begin in section I with a summary of the trajectory theory. In section II, I elaborate, in the context of the trajectory theory, an approach to the problem of preemption that I favor and that has been previously suggested in the context of counterfactual theories of causation but whose versatility I think has not been fully appreciated. This will involve a certain conception of events and a substantive thesis concerning events so conceived. In section III, I discuss the possibility that that thesis might be false, and also discuss the more general phenomenon of *overdetermination* (or *redundancy*, of which preemption, or *asymmetric* overdetermination, is a species).

I. Trajectories

In this section I briefly summarize the probability (or propensity) trajectory theory of singular probabilistic causation.[2] This will be done by noting just five features of the theory. First, it is a theory of *singular* causation. Thus, it seeks to understand such "token level" claims as "Harry's smoking caused his heart attack" or"Suzy's throwing the rock (at that time and place) caused that bottle to break," rather than such "type level" claims as "Smoking is a positive causal factor for heart attacks" or "Throwing rocks at bottles causes bottles to break." The relata of singular causal relations are events that actually occur (understood in a certain way, see below) rather than properties (such as being a smoker or breaking) that can stand in stand in general causal relations. The second point is how events are understood. For the purposes of the theory, they are understood simply as the exemplification of some property or properties at some specified place and time (or interval). Thus, for our purposes, if x stands for a time/place pair (call it $<t_x,s_x>$) and X stands for a property (which may be complex, consisting of conjuncts X_1, X_2, ...), then an event is specifiable by saying that X is exemplified at x; this will be symbolized as "Xx" in what follows. An event is understood simply as the exemplification of a property or properties at such and such a time and place.

Third, the probability trajectory theory is a "probability-increase" theory of singular causation (causes raise the probability of their effects), but not in the ususal sense in which probability increase is understood. In the usual sense, "C increases the probability of E" is understood in terms of conditional probability comparison: Pr(E/C) > Pr(E/~C) (or, equivalently, Pr(C/E) > Pr(E)). With suitable

preempted event is "cut" by some effect of the preempting event at some time intermediate between the times of the preempted event and the effect, where had this "cutting" not occurred the preempted event would have succeeded in causing the effect in question. (See Lewis (1986), and also Lewis (2000) where he prefers the terminology "early and late cutting".) In the Billy and Suzy example, Suzy's rock did not interfere (at least "relevantly much" we might say) with the trajectory of Billy's rock.

[2]For details, see Eells (1991), chapter 6.

qualifications (see below), I think this understanding of probability increase is suitable for understanding probabilistic causation at the type level, that is for understanding claims of the form, for example, "(property) C is a positive causal factor for (property) E." For the token level (for singular causation), however, the probability trajectory theory focuses on the actual evolution of the probability of the token effect from around the time of a candidate token cause event to the time of the token effect event. A probability trajectory is the shape of the time/probability graph that represents such an evolution.

At this point, it will be useful to point out that probabilistic claims must be understood as relative to a population, P, understood to be of a certain kind, Q – where it is the kind Q that really controls, so reference to P will be suppressed in what follows and I will refer simply to "populations Q".[3] For example, certainly the probabilistic (as well as the type level causal) relations between smoking and heart attacks are different in the human (Q) population from what they are in a population of smoking machines (Q') in a laboratory. Thus, type level causal claims should be taken to be of the form "C has such and such a kind (positive, negative, etc., see below) of causal significance for E in population Q" (where the relevant probabilities are understood as relative to Q). At the token level, causal claims will be formulated as "Xx has such and such a kind of token causal significance for Yy," where the relevant population is the unit set {<x,y,>} whose kind can often be taken to be understood by context (as in type level claims) in a request for causal information (in a question of the form "What was the causal role of Xx for Yy?"). Then, with X, x, and y left implicit, I let "$Pr_t(Y)$" symbolize the probability, at time t, that the second member, y, of <x,y> in the singleton population {<x,y>} exhibits (or will exhibit) property Y (the role of X and x in this will be explained shortly). $Pr_t(Y)$ can be thought of as a conditional probability, $Pr(Y/K_t)$, where K_t is the conjunction of relevant (to Y) factors that have fallen into place by time t – but there are two further important features of $Pr_t(Y)$, to which I now turn.

The fourth point about the probability trajectory theory I need to mention is two qualifications analogous to those alluded to above in connection with the type level probabilistic theory of causality. In the type level theory, it will not do simply to say that C is a positive causal factor for E in population Q if and only if $Pr(E/C) > Pr(E/{\sim}C)$, for there is the possibility of "spurious correlation." E can be positively probabilistically correlated with C even when C is not a cause of E, for example if there is a common cause of C and E (for example, rain is positively correlated with falling barometers not because the latter causes the former, or vice versa, but because the two have a common cause, an approaching cold front, or falling barometric pressure). The standard solution to this problem of spurious correlation is to "hold fixed" (conditionalize on) the appropriate factors when making the relevant probability comparisons. Without going into detail, I simply describe the adjustment (for the type level theory, preliminary to the token level theory) to the "basic probability increase idea" (causation goes by correlation) in three steps.[4] First, let $F_1, F_2, ..., F_n$ be those factors that need to be "held fixed" (what they are is given in the third step below), and let $K_1, ..., K_m$ be the maximal conjunctions of the F_i's and their negations that have positive probability both in conjunction with C and in conjunction with ~C (there are $m \leq 2^n$ K_j's). The K_j's are called causal background contexts.

[3]For details, see Eells (1991), chapter 1.

[4]For details, and defense of the requirements, see Eells (1991), chapters 2 and 3.

Second, C is called a "positive", "negative", or "neutral causal factor" for E if and only if, for all j, Pr(E/K_j&C) >, <, or = Pr(E/K_j&~C). The feature of this according to which the inequality or equality must hold across *all* causal background contexts is called *context unanimity.* In case unanimity fails, C is called "causally mixed" for E. And third, we must say what factors should be included among the F_i's. These are all factors that are both (i) causally independent of C (C is causally neutral for them) and also (ii) either (iia) causally relevant (positive, negative, or mixed) for Y or (iib) interactive for Y with respect to X. A factor F is interactive for Y with respect to X if the comparison (greater than, less than, equal to) between Pr(E/F&C) and Pr(E/F&~C) is different from what it is between Pr(E/~F&C) and Pr(E/~F&~C) (i.e., the comparison between Pr(E/C) and Pr(E/~C) is different in the presence of F from what it is in the absence of F).[5]

As to (i) just above, we should not hold fixed factors causally intermediate between C and E in evaluating the causal role of C for E; otherwise, we "rob" C of the probabilistic impact on E that it should have in virtue of its causal impact, if any, on E. (iia) and (iib) just above are the two qualifications alluded to above for the type level theory. (iia) is intended to handle the usual kinds of spurious causation, and (iib) is intended to get the theory to give the right answer of "mixed" causal factorhood when this is the truth.

For our purposes here, the important point (this fourth point about the *trajectory* theory for probabilistic causation on the *token* level) is that qualifications analogous to (iia) and (iib) – subject still to (i) – apply also at the token level. That is, in assessing the token causal significance of an event Xx for an event Yy, we must control for events that occur causally independently of Xx and that bear possibly confounding casual relations (relative to Xx) for Yy. The best way to implement these requirements, I think, is to hold fixed (positively or negatively, depending on how things actually happen) factors Z such that Zz (for the relevant time and place z) actually occurs *token causally independently* of Xx and the factor Z is either (again) either (iia) a positive, negative or mixed cause of Y *at the type level* or (iib) interactive for Y with respect to X *at the type level*. Of course, at the token level, there is just one relevant causal background context, call it K_a, which corresponds to features of the way things actually are in the actual situation in question. Then the relevant probability trajectory traces the evolution of Pr(Y/K_a&K_t) as t varies from around the time of Xx to the time of Yy). As in the type level theory, the background context is a very important feature of the token level theory, and the relevant qualifications on the basic probability increase idea are natural given their analogs in the type level theory.[6]

Fifth, and finally, the taxonomy of kinds of causal significance is somewhat different at the token level, according to the trajectory theory, from the way it is in the type level theory. This is because of the different conception of probability change used at the token level. At the type level, there are,

---

[5]A more general idea of interaction is that the pairs <Pr(E/F&C),Pr(E/F&~C)> and <Pr(E/~F&C),Pr(E/~F&~C)> are different. More general still would be to define interaction in terms of partitions of factors F.

[6]See Eells (1991) for examples, specifically at the token level, that demonstrate the need for holding fixed independent causes and interactive factors.

qualitatively, four kinds of probabilistic impact that C can have on E: unanimously positive, unanimously negative, unanimously neutral, or nonunanimous.(for positive, negative, neutral, and mixed casual factorhood, respectively, and of course there are several ways in which C can be nonunanimous for E, both probabilistically and causally). These are different kinds of conditional probability comparisons across contexts. The token level theory, on the other hand, pays attention to the way the probability of a later event Yy actually evolves from around the time of an earlier event Xx to the time of Yy. And qualitatively speaking, there are four basic shapes such a probability trajectory can assume: 1) it can be higher just after the time of Xx than it was just before that time and stay higher all the way until the time of Yy, 2) it can be lower just after the time of Xx than it was just before that time, 3) it can be the same just after the time of Xx as it was just before that time, and 4) it can be higher just after the time of Xx than it was just before that time but not remain higher than that previous value all the way until the time of Yy. At the time t of Yy, the probability of Y becomes 1 ($Pr(Y/K_a\&K_t) = 1$ when t is the time of Yy or after that time). Note also that value of the probability of Y *at* the time of Xx does not enter into the theory; this I take to be equal to $Pr(Y/K_a\&K_t)$, where t is the time of Xx, and K_t includes X but does not register the (perhaps improbable) actual consequences of the occurrence of X at that time. I have called (just to pick some suggestive terminology) the four kinds of causal significance corresponding to 1)-4) above Yy's happening *because of*, *despite*, *independently of*, and *autonomously of* Xx, respectively. These kinds of causal significance can come in degrees, but I will not enter into that here except to say that the degrees can be measured, basically, by the magnitudes of absolute probability differences for the candidate effect event across the time of the candidate cause event.[7]

I should point out that the explications given above of the various token and type level causal concepts are not intended to be definitions – as such they would be circular, of course. Rather, they should be understood as constraints on the relationships among probabilistic and causal relationships. And finally, I point out that the probability trajectory theory is supposed to apply only in cases of nondeterministic causation – for deterministic causation, the probability trajectories would be trivial and the differences between the four kinds of token causal significance described above could not show up.

II. Probabilistic causal preemption.

For a long time, the phenomenon of preemption, described above, has provided test cases for theories of causation. J. L. Mackie (1974, pp. 44-45) describes several examples, and gives references dating back to the 1920's. And, as mentioned above, the problem of preemption has been used recently to challenge the probability trajectory theory just outlined. I believe that the probability trajectory idea has the resources to deal with the phenomenon – to give the right, indeed the intuitively right, answers about what causes what, when the questions and answers are properly formulated and understood. I begin by giving three examples described by Mackie, and then turn to a couple of more recent examples. The numbering below follows Mackie (1976, p. 44).[8]

---

[7]This will do for present purposes. Again, for details, see Eells (1991), chapter 6.

[8]Quoted from Mackie (1974, p. 44). These examples (iii)-(v) are what today we call "preemption", or "asymmetric overdetermination", where it is clear which of two earlier events is the

(iii) '...conditions (perhaps unusual excitement plus constitutional inadequacies) [are] present at 4.0 p.m. that guarantee a stroke at 4:55 p.m. and consequent death at 5.0 p.m.; but an entirely unrelated heart attack at 4.50 p.m. is still correctly called the cause of death, which, as it happens, does occur at 5.0 p.m.'

(iv) Smith and Jones commit a crime, but if they had not done so the head of the criminal organization would have sent other members to perform it in their stead, and so it would have been committed anyway.

(v) A man sets out on a trip across the desert. He has two enemies. One of them pours a deadly poison in his reserve can of drinking water. The other (not knowing this) makes a hole in the bottom of the can. The poisoned water all leaks out before the traveller needs to resort to this reserve can; the traveller dies of thirst.

Let's focus on case (v) for now. In this case it is supposed to be clear that it is the puncturing of the can, not the poisoning of the water in it, that caused the death, though either alone (or as in the example both together) would have sufficed. Let us see what the probability trajectory theory has to say about this. And for this purpose let us assume that all the causal relations involved are probabilistic and that all the relevant probabilities are nonextreme (not 0 or 1) – that is, until the time of an event, at which time its probability assumes the value 1. First, the theory correctly rules that the poisoning of the water in the can did not cause the death. This is because the puncturing of the can is a cause of death that is causally independent of (not an effect of) the poisoning (the second enemy didn't even know about the poisoning). In fact, all of the relevant effects of puncturing the can, including the presence of terminal dehydration, are causally independent of the poisoning and also causally relevant to the death. Thus, all these factors (including the puncturing and the dehydration) have to be held fixed (conditionalized on) when assessing the probability-trajectory probabilities relevant to assessing the causal role of the poisoning for the death. And given the presence of all these factors, the poisoning does not change the probability of death across the time of the poisoning. So the theory correctly rules that the death is causally independent of the poisoning of the water in the can.

What about the token causal role of the puncturing of the can for death, according to the probability trajectory theory? In this case, we have to hold fixed the poisoning of the water in the can, for this is causally relevant to death and causally independent of (not caused by) the puncturing. There are two things to say here. *First*, focusing on the causal role of the puncturing for the *factor of death*, if the poisoning does not necessitate death, then there is still room for the puncturing to increase (or to decrease, see below) the probability of death, even conditional on the poisoning; but if the poisoning is highly efficacious in producing death (at the type level), then any increase in the probability of death across the time of the puncturing would be very small, but the theory would still rule that the death is

---

cause of the later event. His examples (i) and (ii), to be considered in section III, are cases of what we now call "symmetric overdetermination", where we are supposed to have no definite intuitions about which of the earlier events is the cause. This terminology is due to Lewis (1986). Mackie quotes example (iii) from M. Scriven (1974), (iv) is from K. Marc-Wogau (1962), and example (v) is based on a modification by Hart and Honoré (1959) of an example of J. A. McLaughlin (1925-6).

because of the puncturing, but only to a small degree. Further, however, there is the possibility that, if the poison is more efficacious in producing death than the puncturing of the can is, then the probability of death could actually *decrease* across the time of the puncturing (since then the man becomes no longer vulnerable to the poison but only to the less efficacious cause of death, the puncturing of the can). In that version of the example, the trajectory theory would rule that the death is *despite* (to some small degree I suppose) the puncturing. And if the poisoning and puncturing are equally efficacious for death, then the probability of death could remain the same across the time of the puncturing and the theory will say that the death is causally independent of the puncturing. I think all this (the verdicts of the trajectory theory in the various versions of the example) is correct, when we focus simply on the factor of *death* as what was exemplified by the man at the relevant later time and place. But *second*, if this seems unintuitive (that the degree of causal significance of the puncturing should be called "small" in the example, or that the death should be called even a little "despite", or even "causally independent of", the puncturing), then I think it does so only because it leaves out the rest of the causal story as seen from the point of view of the probability trajectory theory. There are of course many factors that are exemplified at the relevant time and place, including not only death but also the *factor of death-accompanied-by-dehydration*.[9] And of course, even holding fixed the factor of poisoning, we should expect the probability of death-accompanied-by-dehydration to rise considerably across the time of the puncturing, so that the probability trajectory theory will give the correct answer that the death-by-dehydration was, to a significant degree, "because of" the puncturing.

Thus, in example (v), the trajectory theory gives the clearly correct answer about the casual role of the poisoning for death, and when the relevant factors are isolated and the relevant questions asked, the theory gives the correct answers about the causal role of the puncturing for death, and for death-accompanied-by-dehydration. Examples (iii) and (iv) can be handled in analogous ways, when understood to involve nondeterministic causation and nonextreme probabilities.. In (iii), the death is, intuitively and according to the theory, casually independent of those conditions that were right for a stroke (hold fixed the "entirely unrelated" heart attack and the resulting lack of blood circulation that eventually and more proximately led to the death); and the death is, to some small degree, either because of or despite (or even independent of, depending on the details of the example) the heart attack (depending on the relative efficacies of the heart attack and the pre-stroke conditions for death); but death-accompanied-by-lack-of-blood-circulation is, to a high degree, because of the heart attack (hold fixed, of course, the pre-stroke conditions). In (iv), the crime is, intuitively and according to the theory, causally independent of the backup plans of the head of the criminal organization (hold fixed Smith and Jones' intentions and the successes in the various steps along the way that culminated eventually in the crime); and the mere fact of the crime is, to some small degree, either because of or despite (or even independent of, depending on details of the example) Smith and Jones' forming the intention to commit it (depending on the relative skills of Smith and Jones compared to the backup crew); but the crime, in the exact way it was committed by Smith and Jones, is, to a high degree,

[9]Mackie (1976, p. 46) makes much the same point when he distinguishes between the "facts," "[that] *he died*, and that *he died of thirst*."

because of Smith and Jones' forming the intention to commit it (hold fixed, of course, the backup plans of the head of the crime organization). (If however, the plans of the head of the criminal organization in some way contributed to Smith and Jones' making their plans, then the story would be different and we would not hold fixed Smith and Jones' plans in assessing the causal role of boss' plan for the crime, and we would get the right answer that the crime was because of the boss' plans.)

Note that there are three question/answer pairs addressed in the above analysis of the three examples: 1) What is the causal role of the (preempted) event X'x' (poisoning, conditions being right for a stroke, the boss' plans) for effect Yy (death, death, crime), 2) What is the causal role of the (preempting) event Xx (puncturing, heart attack, Smith and Jones' plans) for the effect Yy, and 3) What is the causal role of (preempting) event Xx for the effect event Y'y considered in a more precisely specified way (death accompanied by dehydration, death accompanied by lack of blood circulation, crime in the specific way committed by Smith and Jones). The application of the trajectory theory to questions 1) went the smoothest (by holding fixed the preempting cause and its effects intermediate between it and the final effect). The application to questions 2 was fairly straightforward as well, except that it initially seemed that the trajectory theory was not giving the preempting cause Xx its due, in not assigning it a strong enough causal role in the production of the effect event Yy. However, answers to questions 3) were supposed to fix this seeming lack of match with intuitions by pointing to a more detailed, or different, specification of the way things were at time/place y for which Xx really *was* strongly causally responsible: the idea is that in fact Xx was *not* strongly causally responsible for y's being Y but *was* strongly causally responsible for y's being Y'. Applications of the theory to questions/answers 1) had to do with verifying that *the theory does not say that the preempted cause is a cause*; and applications of the theory to questions/answers 2) and 3) had to do with verifying that *the theory does say that the preempting cause is a cause (of the effect event appropriately understood)*.

While issue 1) seems to be handled just fine by the theory, and the same for 2) given the assumed probabilistic nature of the examples, the application of the theory to issue 3) relies on a substantive assumption, which may be construed either metaphysically or empirically, depending on how one wants to construe the theory. That assumption may be formulated like this:

> *Trace Assumption*: In cases in which Xx preempts X'x' in the production of Yy, there is some feature Y' of what happens at y that physically traces back to Xx and not to X'x' and would not have been present at y had Xx not occurred and X'x' caused Yy instead.

This formulation involves the ideas of "physically tracing back" and "counterfactual dependence"; these ideas are not ingredients in the trajectory theory, but rather this formulation is simply intended to use some (somewhat vague) ideas that are ingredients in other theories of causation and in our ordinary concept of causation and in terms of which we can test the implications of the trajectory theory. In the can of water example, (v), Y' was death specifically accompanied by dehydration; in the patient example, (iii), Y' was death specifically accompanied by lack of blood (ordinarily supplied where and in the manner the heart supplies it); and in the crime example, (iv), Y' had to do with some supposed specific way in which Smith and Jones committed the crime that differs from the way it would have been committed if the backup plan had had to be implemented.

There are three points I would like to make in clarifying and defending this assumption. First, that factor Y' is *not* intended to involve a possible difference in time or place in which Yy did occur as a result of Xx and that would have been different had Xx not occurred and X'x' been the cause of some Y''y' instead. Throughout, I am assuming that we are concerned with the causal significance of what happened at one specific time/place, x, for what happened at another, y. And in fact, if, in the examples, the difference between Xx and X'x''s being the cause did make a difference in the time at which an effect in question occurred, then I think the probability trajectory analysis would have an even easier time dealing with the relevant examples. As explained above, I am working with a conception of events on which they are individuated by a specification of a time/place (point or interval) and a set of factors exemplified at that time/place (like a trope): if either the time/place differed or the set of properties differed (even in one's being a subset of another) in two specifications, then we have specifications of two different events.[10] So if, for example, in the crime case, (iv) above, "the" crime would have been committed at a later time than Smith and Jones' crime had Smith and Jones failed and the backup plan been used, then there would be this temporal feature of the actual crime that traced back to Smith and Jones' plan but not to the boss' backup plan,[11] and surely that there even was a crime at the time/place y of the actual crime is a feature of y that traces back to Smith and Jones' intentions and that would not have been present at y were it not for their plans.

Second, it might be objected that in many cases of preemption there are traces of the preempted event that are present in the effect event that would not be there had the preempted event not occurred, so that my application of the trace assumption does not distinguish the preempted from the preempting event. So, for example, in the patient case, (iii) above, there may be features Y'' present in the patient at the time of his death that physically trace back to the (preempted) conditions-just-right-for-a-stroke-at-4:55-p.m.-and-death-at-5:00-p.m. present at 4:00 p.m., where Y'' would not have been present at y had these 4:00 p.m. conditions not been present. This seems natural enough, of course. And it is also natural enough to say that Y''y *is* because of X'x' but not because of Xx, while Y'y (Y' again = death with lack of blood in the places where the heart ordinarily supplies it in the way it ordinarily does) is because of Xx but not because of X'x', which is just what the probability trajectory theory yields.

And third, we may wonder whether the trace assumption is always true, whether there really could not be cases in which everything that happens at y is just the way it would have been had Xx not occurred and X'x' had been the cause of y's being Y and y's having all the other features Y' it actually has. I know of four examples that have been described recently with this issue specifically in mind. I will consider two of these here.[12]

---

[10]Compare Jaegwon Kim (e.g., 1973), though he individuates events by triples consisting of an individual (or individuals), a property (or properties), and a time.

[11]See L. A. Paul (1998, 2000) for an application of this feature of at least some cases of preemption to the counterfactual analysis of causation.

[12]The other two, by Schaffer (2000), involve either an admittedly unrealistic but logically possible situation involving magic and action at a distance or a possible world with a physics different

L. A. Paul (2000) describes an example involving two cats and a fly:

> C. Louise crouches aiming for [a] fly. Possum also crouches, aiming for the same fly. C. Louise jumps. Possum, who has been practicing, jumps a moment later, but his (newly acquired) agility makes him able to catch the fly at the same time as C. Louise. Unfortunately for Possum, there is a little known law that states that flies, when pounced upon by multiple cats, are captured by the cat who jumps first. Since C. Louise jumps before Possum, she gets the fly. If C. Louise had not jumped, Possum would have captured the fly in the very same way and at the very same time. C. Louise's pounce, albeit through no intrinsic feline merit, trumps Possum's. (Paul 2000, p. 247)

But surely there are features of the actual fly-catching that trace back to C. Louise's pounce and not to Possum's and that would not have been present had C. Louise not pounced and Possum had caught the fly. For example, it was C. Louise's paw (or teeth, or however the cat did it) that actually made contact with the fly; this clearly traces back to C. Louise's pounce, not Possum's, and would not have been present had C. Louise not pounced and Possum had caught the fly instead.[13]

Ned Hall (2001, p. ///) says, "it's easy enough to construct cases in which **c** is clearly a cause of **e** but in which neither **c** nor any event causally intermediate between it and **e** make the slightest difference to the way **e** occurs," and he attributes to Steve Yablo the following modification of the story of Billy and Suzy (described at the beginning of this paper):

> This time Billy throws a Smart Rock, equipped with an on-board computer, exquisitely designed sensors, a lightning-fast propulsion system – and instructions to make sure that the bottle shatters in exactly the way it does, at exactly the time it does. In fact, the Smart Rock doesn't need to intervene, since Suzy's throw is just right. But had it been any different – indeed, had her rock's trajectory differed in the slightest, at any point – the Smart Rock would have swooped in to make sure the job was done properly. (Hall 2001, p ///)

But again, I think that when we think carefully about what happened at y (the time and place of the bottle's shattering) we will find there features that trace back to Xx (Suzy's throwing her rock) that would not have been present had she not thrown her rock, or if for some other reason the Smart Rock had done the job instead. For example, Suzy's rock is in front of the Smart Rock – closer to the bottle than the Smart Rock is just before the shattering, and in contact with the bottle, while the Smart Rock is not, at the time of the initial rock-bottle contact. And just after the initial rock-bottle contact, surely the configuration of Suzy's rock, the Smart Rock, and the shards of glass is different from the way it would

---

from actual physics. While it is interesting to apply our concept of cause to such conceptually possible scenarios, I will restrict my attention here to what seem to be attempts to describe more realistic such situations. The first example to be described would seem to be a more realistic example of the kind that Schaffer has in mind and is intended to illustrate the possibility of "trumping," as Schaffer calls it.

[13]Paul points out that this example involves action at a distance and suggests that a different analysis is necessary for such causation than for causation that needs chains of events. In a note below, I will describe a version of the C. Louise/Possum case in which I think the trace assumption fails and there is no action at a distance (but which is not a case of *late* preemption).

have been had Suzy not thrown her rock, or if for some other reason the Smart Rock had done the job instead. So, realistically and carefully understood, it seems that there are features Y' just before, at the time of, and just after the shattering that trace back to Suzy's throw and that would not have been present had Suzy not thrown her rock, or if for some other reason the Smart Rock had done the job instead. And the probability of these features Y' of y would seem to increase across the time at which Suzy actually succeeds in letting her rock fly. In section III, however, I will discuss the theoretical possibility of examples in which there are no such factors Y', that trace back to Xx and that would not have been present were Xx not to have occurred.

The discussion here may remind readers of David Lewis' discussion of *fragility* of events (1986, p. 196ff.). An event is "*fragile* if, or to the extent that, it could not have occurred at a different time, or in a different manner. A fragile event has a rich essence; it has stringent conditions of occurrence." (p. 196) The idea I think is that what happens at a time and place y can be considered in a more or less detailed way, Y, so that at that same time and place both more and less fragile events occur. ("Don't say: here we have *the events* – how fragile are *they*....Properly posed, the question need not have a fully determinate answer, settled once and for all. Our standards of fragility might be both vague and shifty." (pp. 196-197)) And Lewis considers a strategy like the one above involving our trace assumption, and asks, "Wouldn't we still have residual cases of redundancy [of which preemption is one kind], in which it makes *absolutely* no difference to the effect whether both of the redundant causes occur or only one?", and answers, "Maybe so; but probably those residual cases would be mere possibilities, far-fetched and contrary to the ways of this world." (p. 197, Lewis' italics) The trace assumption is that such far-fetched cases just do not happen (but, again, I will consider such theoretically possible cases in section III). But Lewis goes on to say that the strategy makes for more trouble than it cures anyway, and he considers two examples intended to illustrate this, to which I now briefly turn.

In the first example, one gentle soldier in an eight-soldier firing squad did not shoot. If the victim's death is considered to be a very fragile event, then a seven-bullets-through-the-heart death is a different event from an eight-bullets-through-the-heart death. "So the gentle soldier caused the death by *not* shooting, quite as much as you caused it by shooting! This is a *reductio*." (p. 198) In the other example, Boddie eats first a large dinner and then poisoned chocolates. Boddie then dies from the poison, but the large dinner slowed the absorption of the poison in the chocolates and the death occurred somewhat later and in a slightly different manner than it would have without the large dinner. "If the death is extremely fragile, then one of its causes is the eating of the dinner. Not so." (p. 198) Leaving aside the factor of time (as before), the number of bullets through the heart and the exact manner of Boddie's death are factors Y' that trace back to the gentle soldier's not shooting and the eating the large dinner, respectively, and which would not have been exemplified at y (times/places of the deaths) had the failure to shoot or the eating of the large dinner not occurred.

In these two examples, as in the other examples considered above, there are various things going on in, or features of, the spatio-temporal region y of the effect events, Yy, Y'y, and so on, as well as various things going on in, or features of, the spatio-temporal region x of the cause events, Xx, X'x, and so on. And different features X of the earlier time/place x have different causal significances for different features Y of the later time/place y. While it does indeed sound unusual to say that the gentle soldier's not shooting caused *the victim's death*, or that Boddie's large dinner caused *Boddie's death*, I suggest that this is only because the italicized phrases here single out or bring to mind a feature of the later events – namely Y = *death* – for which the earlier event was not causally responsible. But, as in the previous examples, there are also features – Y' = death-with-exactly-seven-bullets-through-a-heart and Y' = death-by-poison-mixed-with-that-large-dinner – for which the cause events Xx were responsible. I find it completely natural to say that the gentle soldier's omission was responsible for there being only seven bullets penetrating the victim's heart and that Boddie's large dinner caused his death in the way that the death in fact occurred – all under the circumstances (the other seven soldiers' shooting and Boddie about to eat the poisoned chocolate). I agree that fragility is often a matter of vagueness or shiftiness, but when the relevant features of a cause event and an effect event are settled either by context (and the relevant features can shift from context to context) or by asking a precise question (of the form "What was the causal significance of x's being X for y's being Y?"), then a precise answer (which specifies each of x, X, y, and Y) should not, I think, sound unnatural at all.[14]

Before leaving the topic of ordinary preemption (or cases in which it seems, to me at least, that a plausible candidate for a Y' of y tracing back specifically to Xx can be found), I consider two recent examples specifically addressed to probabilistic causation and with the trajectory theory in mind. Christopher Hitchcock describes the following scenario (similar to the Billy and Suzy example described above but not exactly a case of preemption):

> Suppose that two gunmen are shooting at a Ming vase. Each one has a fifty percent chance of hitting the vase, and each one shoots independently, so the probability that the vase shatters is .75. (For simplicity, we will ignore the possibility that the vase might survive a bullet hit.) As it happens, the first gunman's shot hits the vase, but the second gunman misses. (2001, p. ///)

In this case, in evaluating the causal role of either gunman's shooting for the shattering of the vase, we must hold fixed, for one thing, the other gunman's shooting, for the latter is causally independent of the former and causally relevant to the vase's shattering. This gives the right answer for the first gunman's shooting: the probability of the vase's shooting increases from .5 to .75 across the time of his shooting (and I suppose doesn't decrease between that time and the time of the vase's shattering). But at first it may seem that the same is true for the *second* gunman's shooting, and that the trajectory theory would give the wrong answer that the vase's shattering was because of *this* shooting. (As before, I ignore the time element, assuming that the time of the shattering would be the same no matter which gunman's bullet hit the vase.)

However, there are two ways in which the trajectory theory blocks this conclusion. First, it is after all *the first gunman's bullet* that hits the vase; this is a feature of what happens at the place/time y of the shattering that traces back to the first gunman's shooting and that would not have been present at y had the first gunman not shot, and whose presence at y the theory says is because of first gunman's shooting, and of course *not* because of the second's. This is similar to what I said about the version of the story of Billy and Suzy with Billy's Smart Rock. And second, recall that we must hold fixed *all* factors whose exemplifications are token causally independent of a candidate cause and which are

[14] Compare Eells (1991, pp. 286-289).

causally relevant to the relevant factor exemplified at the later time. In this case, the first gunman's bullet entering or being just about to enter the vase (just before the vase actually shatters) is such a factor. And conditional on this factor, the second gunman's shooting does not affect the probability of the vase's shattering across the time of his shooting, which makes the theory give the correct answer that the shattering was token causally independent of the second's gunman's shooting.[15] In addition, presumably at some time after the two shots, the bullet from the second gunman is off course; this should be held fixed in evaluating the causal role of the first gunman's shooting for the shattering, in which case the probability of the shattering increases from 0 to .5 at the time of the first gunman's shooting. This would seem to be the correct way to apply the theory to the question of the casual significance of the first gunman's shooting for the vase's shattering, in which case again the theory says that the shattering was because of the first gunman's shooting.

Douglas Ehring (1994, 1997) describes a somewhat more complicated example. A very sick patient in a hospital is connected to a mechanism that can deliver drugs to the patient. The mechanism is connected to two sources of drugs, a source of drug A and a source of drug B. If the patient receives neither of these drugs, then the chances of survival are very low; but if the patient gets either or both of the drugs, then the chances of survival are .5; and there is a law of biology according to which, given the patient's condition, the chances of survival cannot possibly exceed .5 at this time. It is improbable that both drugs would flow to the mechanism that is connected to the patient. And if both drugs are released from their sources, then there is a chance that the drug that reaches the mechanism first will set off a device that blocks the flow of the other drug to the patient. The release of either drug from its source is causally independent of whether the other drug is released from its source. And assume finally that if at least one of the drugs is released from its source then the probability is 1 that the patient will receive one of the two drugs through the mechanism (drug A if only A is released from its source, B if only B is released from its source, and at least one of the two if both are released from their sources). Here is what happens. First, a valve is opened that starts the flow of drug B to the mechanism; then (improbably) a valve is opened that starts the flow of drug A to the mechanism. But drug A reaches the mechanism first and this (improbably) blocks drug B from flowing through the mechanism to the patient; drug A reaches the patient and the patient survives. Ehring says (correctly, in a way) that the release of drug A is what actually saved the patient. However, at the time of A's release, the probability that the patient would survive was already at its maximum value of .5, since drug B was already on its way to the mechanism, guaranteeing that at least one of the two drugs would reach the patient. Thus, the release of drug A did not, at the time of its occurrence, alter the probability that the patient would survive. This is supposed to be a counterexample to the probability trajectory theory since the release of drug A caused the survival but the probability trajectory theory says that survival was causally independent of the release of A.

However, I think the right way to view this example is this: the release of drug A *did not* cause (exactly) *the survival*, but it did cause *the mechanism to deliver the drug it did (A)* to the patient, it

---

[15]Hitchcock in fact makes this point for what he calls "ebb and flow in the probability pool" approaches to singular probabilistic causation, of which the probability trajectory theory is an instance.

did cause the mechanism to start delivering the drug it did *at the time it did (just before drug B reached the mechanism)*, it did cause *the mechanism to block the flow of drug B*, and it did cause the patient to *survive with drug A in his veins*. Even though the probability of survival cannot be increased above the stipulated biological limit of .5 by the release of drug A, the probabilities of these other factors can be increased by this from what their values were before the release of drug A, but of course not of the last other factor to above .5. Again, if we pay attention to the fact that many properties are exemplified at the relevant later time/place y (in this example, the spatio-temporal region within which the patient recovered), then, when we tell the whole story, the probability trajectory theory will deliver the correct answers, in the light of which some initially perhaps misleading answers (in this example, that the release of drug A *did not* in the circumstances cause, exactly, the patient to survive) are no longer misleading. (I leave it to the reader to ponder the application of the probability trajectory theory to the causal role of the release of drug B for the exemplifications of the various factors exemplified around the patient's recovery.)

I conclude that in cases in which the trace assumption holds, and when causal questions and answers are formulated precisely (including an x, an X, a y, and a Y), the trajectory theory gives the correct answers – indeed the *intuitively* correct answers in light of answers to other questions, that is, in light of the "whole story."[16]

III. Failure of tracing, and symmetric overdetermination

In this section I discuss two matters that were alluded to above but postponed, namely the theoretical possibility of failure of the trace assumption and the phenomenon of symmetric overdetermination.

Theoretical possibility of failure of the trace assumption. Suppose that, somehow, an effect e would have occurred in exactly the same way if, instead of the preempting cause c causing e, the preempted cause c' had caused e. Expressed in terms of the notation introduced above, the situation is this: for *all* factors Y present at time/place y, y would still have exemplified Y had X'x' not been preempted and Xx was not causally responsible for what went on at y. As suggested in the discussion of examples above, I think this would be extremely rare in a case of (at least late) preemption,[17] so I will turn to a different kind of case, an example basically the same as one I discussed

[16] I note also the uniformity of this treatment of preemption across cases of "early" and "late" preemption. The versions of the Billy and Suzy story discussed are clear cases of late preemption, the case of the thirsty traveler and Ehring's patient example are clearly cases of early preemption, and in some of the other examples discussed, different verdicts concerning classification into late or early preemption may depend on different understandings or specifications of the details of the examples.

[17] But consider this modification of the C. Louise/Possum case of Paul's described above. Suppose C. Louise sees Possum on the way to the fly and pushes Possum away and is able then to accelerate so that the time of the effect would be the same no matter which cat got the fly. Further, more saliently, suppose the effect is not exactly the catching of the fly but rather the fly being crushed inside of a box by a mechanism that is activated outside the box by a button, where the exact way the

briefly in my (1991, pp. 384-386) that I think will serve our purposes here but which I am not sure should be counted as a case of preemption.[18]

Suppose a golf ball is rolling straight toward the cup with a 95% chance of falling in. A squirrel kicks the ball away, but improbably enough the ball comes off the squirrel's foot with the same 95% chance of falling into the cup, but along a different path (made possible by the contour of the golf course). Suppose further that this new 95% chance is for the ball's landing in the cup in exactly the same way ("locally") as it had a 95% chance for before the kick (i.e., crossing the same point on the rim of the cup, from the same direction, with the same speed, at the same time, etc.). The ball falls in the cup.[19] In one sense, we want to say that the squirrel kick caused the birdie, but on the other hand it is part of the example that the actual effect of the kick is to leave the probability of the birdie's occurring, in exactly the same way as it probably otherwise would have, unchanged. Even though the birdie's happening exactly the way it did traces back to the squirrel kick, the net actual effect of the kick is to leave the probability of this happening unchanged. Because of the latter, the trajectory theory says the birdie occurred (exactly as it did) "token causally independently" of the kick. This seems to be a correct verdict of the trajectory theory in one sense of cause: the probability of things being the way they actually turned out to be at the time of the birdie was left unchanged by the squirrel kick, given the way the ball was moving after the squirrel kick. But there is still the tracing-back intuition according to which the squirrel kick *is* relevant to the birdie, even though there is no feature of the effect that traces back to the squirrel kick and that would (with the same probability as conferred by the ball's original motion) not have been present had the squirrel not kicked the ball..

My intuition here is that, as far as causation strictly speaking in concerned, the squirrel kick *is* irrelevant to the birdie's happening, in the exact way that it did – that being, again, because the kick left the probability of the ball's falling into the cup, in exactly the way it did, unchanged. On the other hand, there is the "tracing-back" intuition that the squirrel kick is somehow relevant to the birdie, for the kick was responsible for ball's taking the path it actually did take into the cup. If we want to explain just how or why the ball fell into the cup, we would surely want to include the fact of the kick and the fact of the path the ball actually took in moving to the cup. The way I have just expressed the way in which I

---

fly is crushed is not affected by who or how the button is pushed. Of course this is not a case of *late* preemption (which was Paul's topic) but it would seem to be a case of failure of the trace assumption (and in which there is no action at a distance). So I do not deny the possibility of failure of the trace assumption in cases of (at least *early*) preemption.

[18]I will note below why I think it is unclear that this should be counted as a case of preemption.

[19]Why isn't this clearly a case of preemption, of the golfer's swing or of the motion of the ball just before the squirrel's kick, by the squirrel kick? Each of these candidates for the preempted cause seems clearly part of the actual cause of the birdie: on a natural understanding of the example, had the golfer not swung, or the ball not been moving in the way it was just before the kick, the action of the squirrel would not have resulted in a birdie. So the influences of these candidates for a preempted cause are clearly not totally preempted by the action of the squirrel.

think the kick is relevant to the birdie suggests my diagnosis of the case: we should separate the concepts of causation and explanation here. It would be a project in itself to develop this suggestion in detail. I suggest that we should separate question of the explanatory contribution of the kick for the birdie (in the exact way that it happened, in this unusual case) from the question of causal impact of the kick on the birdie (the kick's having no actual effect on the probability of the birdie in the exact way that it happened, in this unusual case).

I think this intuition and suggestion applies also to the theoretically possible kind of preemption in which the trace assumption fails. In such cases (if there are any), y's being Y (for all relevant Y's exemplified by y) traces back to and is partially *explained* by Xx, even though, due to the presence of the backup event X'x', Xx did not, in these circumstances, positively *causally* contribute to y's being Y. I think that this tracing back, explanatory kind of significance of an earlier event for a later event can be captured by a trajectory-style analysis. It is not clear to me exactly how this should be carried out, but it seems that the basic idea should be to point to chains or networks of trajectory-style connections between an earlier event and a later event. For example, while the birdie was not exactly because of the squirrel kick, there are intermediate events that were because of the kick and that the birdie was because of (thus, "because of" is not in general transitive[20]). But again, this is a topic for a separate project.

Symmetric Overdetermination. I turn now to the second kind of example of overdetermination, exemplified in examples (i) and (ii) in Mackie's list of five examples, (iii)-(v) of which were discussed above.

> (i) A man is shot dead by a firing squad, at least two bullets entering his heart at once, either of which would have been immediately fatal.
>
> (ii) Lightning strikes a barn in which straw is stored, and a tramp throws a burning cigarette butt into the straw at the same place and at the same time: the straw catches fire. (Mackie 1974, p. 44)

Mackie says, "In these cases even a detailed causal story fails to discriminate between the rival candidates for the role of cause. We cannot say that one rather than the other was necessary in the circumstances even for the effect *as it occurs*." (p. 47, Mackie's italics) Of course the concern in this paper is with probabilistic causation, so let us understand the examples as probabilistic (what the members of the firing squad did makes the victim's death very probable and the lightning and the tramp's cigarette butt each make the fire very probable), so that the necessity referred to in the quote from Mackie is not exactly to the point in the way we will understand the examples. The salient point about these examples is indicated by Mackie's italicized phrase. Unlike Mackie's other examples (where we can say that the effect event was a heart-attack-death and not a stroke-death, the effect was a Smith-and-Jones-crime and not a back-up-crew-crime, and the effect was a dehydration-death and

---

[20]Hall (2001) distinguishes between two concepts of cause which he calls "dependence" (which is just counterfactual dependence) and "production" (which is a tracing-back idea). He gives (deterministic) examples of dependence without production; the example just discussed can be thought of as a (probabilistic) case of production without dependence.

not a poison-death), examples (i) and (ii) are supposed to be cases in which the effect cannot be traced back to one of the two candidate causes and not the other.

I think the same considerations applied above to cases of preemption, or asymmetric overdetermination, can be applied also to examples like (i) and (ii), and that the probability trajectory theory, properly applied, yields intuitively correct verdicts. First, although it is not specified in the examples, it seems natural to think that, understanding the examples probabilistically, the combination of two causes makes for a higher probability of the effect than one of the two alone does (of course it is natural for this unspecified possible feature of the examples not be specified by Mackie, since he was operating with a necessity/sufficiency conception of causation). Also, in evaluating the causal role of one of the earlier events for the later one we should hold fixed the other earlier event (since it is at least implicitly assumed in the examples that the two earlier events are causally independent of each other). In this case, we should say that all of the earlier events were causes of the respective later events in their examples. And the probability trajectory theory applied to these examples says that the victim's death was because of each of the bullets' entering the heart (each raised the probability of death over what it is conditional only on the other) and the fire was because of the lightning and also because of the cigarette butt (each raised the probability of the fire over what it is conditional only on the other). However, the examples could be understood in a way that, given that one of the earlier events occurs, the other cannot further increase the probability of the relevant later event (as in Ehring's example in which there is no possibility of raising the probability of the patient's survival above .5). Also, as in previous examples, one could naturally complain that this approach does not assign an intuitively high enough causal significance to the causes for the effect.

But second, we could try applying the trace assumption to these examples. Mackie says that the two candidates for a cause cannot be discriminated between even by considering the effect "*as it occurs*," by which I take him to mean, "paying attention to all of the features of the later event." Of course it is worth pondering this theoretical possibility, but again I find it hard to imagine that the effect event would be exactly the same, in all its details, if either one of the two candidate causes occurred without the other. Surely there is some feature Y' of the firing squad's victim's death that would have been different had only one bullet pierced his heart (e.g., a difference in the change of momentum undergone by a part of the victim's body) and that traces back to a second bullet. And in the fire example, surely the presence of the lit end of the cigarette butt (or of the lightning) made for a difference Y' in what things were like at the time and place y of the fire's ignition that would not have been present if only the lightning (or lit end of the cigarette butt) had been there to start the fire, and which traces back to the presence of the cigarette butt (the lightning strike). And these differences Y' can be expected to reverberate into the exact future of the victim's body and into the exact ensuing course of the fire. In this case, we can naturally say – and the probability trajectory theory applied to the examples yields these verdicts – that the feature Y, the firing squad's victim's being dead, of the time and place of what happened to the victim is token causally independent of each of the bullet's entering the heart (under the circumstances of the other bullet's entering the heart), while the feature Y', above, of this time and place was because of a given bullet's entering the heart, and similarly for the fire example.

But third, we should not ignore the theoretical possibility of failure of the trace assumption in examples like this. In this case, I would again, as above, appeal to a distinction between a purely causal concept and an explanatory concept. If it is really true that neither of the "candidates for the role of cause" made any difference at all in the effect event (under these circumstances of the presence of the other candidate) then it seems, at least to me, reasonable to say that neither *was* a *cause* (under the circumstances), though a full *explanation* of just how and why the victim died or of how and why the fire was started would include both "candidates."

IV. Conclusion

The propensity trajectory theory of singular causation provides a natural framework which supports intuitive judgments concerning causal roles of preempting, preempted, and symmetrically overdetermining evens. (And the phenomenon of preemption poses no special problem for the propensity trajectory theory of singular probabilistic causation.) Some of the conclusions above depend on the trace assumption, and given the theoretical possibility of failure of the trace assumption, on a suggested way of separating purely causal from otherwise explanatory roles of events for events.

REFERENCES

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

Ehring, Douglas (1994) "Preemption and Eells on Token Causation," *Philosophical Studies*, Vol. 74, pp. 39-50.

Ehring, Douglas (1997) *Causation and Persistence: A Theory of Causation*. Oxford University Press, New York and Oxford.

Hall, Ned (2001) "Two Concepts of Causation," *The Journal of Philosophy*, forthcoming.

Hart, H. L. A. and Honoré A M. (1959) *Causation in the Law*. Oxford University Press, Oxford.

Hitchcock, Christopher (2001) "Do All and Only Causes Raise the Probabilities of Effects," in John Collins, Ned Hall, and Laurie Paul (eds.), *Causation and Counterfactuals*, M. I. T. Press, Cambridge, Massachusetts and London, England, forthcoming.

Kim, Jaegwon (1973) "Causation, Nomic Subsumption, and the Concept of Event," *The Journal of Philosophy*, Vol. LXX, No. 8, April 26, 1973, pp. 217-236.

Lewis, David (1973) "Causation," reprinted in his *Philosophical Papers, Volume II*. 1986, Oxford University Press, New York and Oxford.

Lewis, David (1986) "Postscripts to 'Causation'," in his *Philosophical Papers, Volume II*. Oxford University Press, New York and Oxford.

Lewis, David (2000) "Causation as Influence," *The Journal of Philosophy*, vol. XCVII, No. 4, April 2000, pp. 182-197.

Mackie, J. L. (1974) *The Cement of the Universe: A Study of Causation*. Oxford University Press, Oxford.

Marc-Wogau, K. (1962) "On Historical Explanation." *Theoria*, Vol. xxviii, pp. 213-233.

Menzies, Peter (1996) "Probabilistic Causation and the Pre-emption Problem," *Mind*, Vol. 105, No. 5, pp. 85-117.

McLaughlin, J. A. (1925-6) "Proximate Cause," *Harvard Law Review*, Vol. xxxix, pp. 149ff.

Paul, L. A. (1998) "Keeping Track of Time: Emending the Counterfactual Analysis of Causation," *Analysis*, Vol. LVIII, No. 3, pp. 191-198.

Paul, L. A. (2000) "Aspect Causation," *The Journal of Philosophy*, Vol. XCVII, No. 4, April 2000, pp. 235-256.

Schaffer, Jonathan (2000) "Trumping Preemption," *The Journal of Philosophy*, Vol. XCVII, No. 4, April 2000, pp. 165-181.

Scriven, M. (1964) "Review of *The Structure of Science* by E. Nagel," *Review of Metaphysics*, Vol. xviii, pp. 403-424.