

Metaphysics over Methodology--Or, Why Infidelity Provides No Grounds To Divorce Causes from Probabilities

[from M-C Galavotti, P Suppes, and D Costantini (eds) *Stochastic Causality* 2001]

David Papineau

1 Introduction

A reduction of causation to probabilities would be a great achievement, if it were possible. In this paper I want to defend this reductionist ambition against some recent criticisms from Gurok Irzik (1996) and Dan Hausman (1998). In particular, I want to show that the reductionist programme can be absolved of a vice which is widely thought to disable it--the vice of infidelity.

This paper also carries a general moral. It is dangerous to muddle up metaphysics with methodology. If you are interested only in the methodological question of how to find out about causes, you will be unmoved by my defence of reductionism, since it hinges on metaphysical matters that are of no methodological consequence. Indeed, if you are interested only in methodological matters, you may as well stop reading here, since my reductionism will offer no methodological improvement over the non-reductionist alternatives.

On the other hand, if you are interested in the underlying structure of the universe--and in particular in how there can be causal direction in a world whose fundamental laws are symmetrical in time--then I may have something for you. I admit my favoured theory offers nothing new to market researchers who want to find out whether some form of advertising causes improved sales. But I can live with that, if my theory explains the arrow of time.

For any readers new to this area, I should explain that the kind of reductive theory at issue here only has partial reductive ambitions, in that it takes probabilistic laws as given, and then tries to explain causal laws on that basis. The hope is to weave the undirected threads of probabilistic law into the directed relation of causation. Perhaps it would be more helpful to speak of a reduction of causal direction, rather than of causation itself. If we think of causal laws as being built from two components--first a symmetrical lawlike connection linking effect and cause, and second a causal 'arrow' from cause to effect--then the reductive programme at issue here aims only to reduce this second directional component. In particular, it aims to reduce it to facts involving the first kind of component (more specifically, to undirected probabilistic laws between the cause, effect and other events). However, it does not aim to explain these probabilistic laws themselves. (It says nothing, for example, about the difference between laws and accidental frequencies).

Most recent discussion of the reductionist programme has focused on methodological rather than metaphysical issues (including Papineau 1993a). This is understandable, to the extent that it is the possibility of real-life inferences from correlations to causes which motivates most technical work in this area. My strategy in this paper will be to place this methodological work in a larger metaphysical context. This metaphysical context won't make any difference to the methodology, but, as I said, methodological insignificance seems a small price to pay, if we can explain why causation has a direction.

2 Probabilistic Causation and Survey Research

A good way to introduce the metaphysical issues will be to say something about 'probabilistic causation' generally. Over the past three or four decades it has become commonplace to view causation as probabilistic. Nowadays, the paradigm of a causal connection is not C determining E, but C increasing the probability of E-- $P(E/C) > P(E)$. (Suppes, 1970.)

This shift in attitudes to causation is often associated with the quantum mechanical revolution. Given that quantum mechanics has shown the world to be fundamentally chancy, so the thought goes, we must reject the old idea of deterministic causation, and recognize that causes only fix quantum mechanical chances for their effects.

However, there is another way of understanding probabilistic causation, which owes nothing to quantum metaphysics. This is to hold that probabilistic cause-effect relationships arise because our knowledge of causes is *incomplete*. Suppose that C does not itself determine E, but that C in conjunction with X does. Then $P(E/C)$ can be less than one, not because E is not determined, but simply because X does not occur whenever C does.

A useful label for this possibility is 'pseudo-indeterminism' (Spirtes, Glymour and Scheines, 1993). If we focus only on C, and ignore X, it will seem as if E is undetermined. But from a perspective which includes X as well, this indeterminism turns out to be illusory.

This pseudo-indeterministic perspective in fact fits much better with intuitive thinking about 'probabilistic causation' than quantum metaphysics. The real reason contemporary intuition associates causes with probabilities is nothing to do with quantum mechanics. (Indeed, when we do look at real microscopic quantum connections, causal ideas tend to break down, in ways I shall touch on later.) Rather, all our intuitively familiar connections between probabilistic and causal ideas have their source in survey research, or less formal versions of such research--and to make sense of these research techniques, we need something along the lines of pseudo-indeterminism, not quantum mechanics.

Let me explain. By 'survey research' I mean the enterprise of using statistical correlations between macroscopic event types to help establish causal conclusions. To take a simple example, suppose that good exam results (A) are correlated with private schools (B). Then this is a prima indication that schools exert a causal influence on exam results. But now suppose that in fact private schools and good exam results are correlated only because both are effects of parental income (C). If that is so, then we would expect the school-exam correlation to disappear when we 'control' for parental income: among children of rich parents, those from state schools will do just as well in the exams as those from private schools; and similarly among the children of poor parents.

In this kind of probabilistic case, C is said to 'screen off' A from B. Once we know about C (parental income), then knowledge of B (school type) no longer helps to

predict A (exam results). Formally, we find that the initial correlation-- $P(A/B) > P(A)$ --disappears when we condition on the presence and absence of C: $P(A/B \& C) = P(A/C)$ and $P(A/B \& \neg C) = P(A/\neg C)$.

To continue with this example for a moment, focus now on the correlation between parental income and exam results itself. Suppose that survey research fails to uncover anything which screens off this correlation, as parental income itself screened off the initial correlation between schools and exam results. Then we might on this basis conclude that parental income is a genuine cause of exam results.

Inferences like these are commonplace, not just in educational sociology, but also in econometrics, market research, epidemiology, and the many other subjects which need to tease causal facts out of the frequencies with which different things are found together.^[1] Now, it is a large issue, central to this paper, whether any causal conclusions ever follow from such statistical correlations alone, or whether, as most commentators think, statistical correlations can only deliver new causal facts if initially primed with some old ones ('no causes in, no causes out'). But we can put this issue to one side for the moment. Whether or not survey research requires some initial causal input before it can deliver further causal output, the important point for present purposes is that, when survey research does deliver such further conclusions, these conclusions never represent purely chance connections between cause and effect.

Suppose, as above, that survey research leads to the conclusion that parental income is a genuine cause of exam results. Now, the soundness of this inference clearly doesn't require that nothing else makes a difference to exam results, apart from parental income. For parental income on its own clearly won't fix a pure chance for exam results. Other factors, such as the child's composure in the exam, or whether it slept well the night before, will clearly also make a difference. All that will have been established is that parental income is one of the factors that matters to exam results, not that it is the only one. As it is sometimes put, parental income will constitute an 'inhomogeneous reference class' for exam results, in the sense that different children with the same parental income will still have different chances of given exam results, depending on the presence or absence of other factors. ($P(E/C \text{ and } X) \neq P(E/C \text{ and } \neg X)$.)

This point is often obscured by worries about 'spurious' correlations. If we want to infer, from some initial correlation between C and E, that C causes E, we do at least need to ensure that C rather than not C still increases the probability of E when we

^[1] Such research is usefully thought of as proceeding in two stages. First, we need to get from finite sample data to lawlike probability distributions; second, we need to get from these probability distributions to causal structure. The first stage hinges on standard techniques of statistical inference. In this paper I shall say nothing about the logic of such techniques, important as this subject is, and simply assume knowledge of lawlike probability distributions. My focus here is exclusively on the second stage, which takes us from lawlike probabilities to causal conclusions. This division of labour is in line with the limited ambitions of my reductive agenda: as I explained at the beginning, the aim is only to reduce causal direction to lawlike probabilistic connections, not to reduce the latter in turn.

'control' for further possible common causes. ($P(E/C \text{ and } X) > P(E/\text{not-}C \text{ and } X)$ and/or $P(E/C \text{ and not-}X) > P(E/\text{not-}C \text{ and not-}X)$.) The point of survey research is precisely to check, for example, whether or not the parental income-exam correlation can be accounted for by the spurious action of some common cause. Thus in practice we need to check through all possible common causes of C and E, and make sure that C still makes a difference to E after these are held fixed. This might make you think that survey research needs to deal in pure homogeneous chances after all. For haven't I just admitted that we are only in a position to say C causes E when we know the probability of E given C and $X_1 \dots X_n$, when these Xs are all the other things which make a probabilistic difference to E?

No. I said we need to check for all possible common causes of C and E. I didn't say we need to check through all other causes of E tout court. This difference is central to the logic of survey research. Before we can infer a cause from a correlation, we do indeed need to see what difference any common causes make to the probability of E. But we don't need to know about every influence on E. This is because most such influences will be incapable of inducing a spurious correlation between C and E. In particular, this will be true whenever these other influences are themselves probabilistically independent of the putative cause C. If some other cause X (good night's sleep) is probabilistically independent of C (parental income), then it can't generate any spurious C-E correlation: X will make E more likely, but this won't induce any co-variation between C and E, given that C itself doesn't vary with X. So survey research can happily ignore any further causes which are probabilistically independent of the cause C under study. The worrisome cases are only those where the further cause X is itself correlated with C, since this will make C vary with E, even though it doesn't cause E, because it varies with X, which does.

The moral is that you don't need to gather statistics for every possible causal influence on E whenever you want to use survey data to help decide whether C causes E. You can perfectly well ignore all those further influences on E (all those 'error terms') that are probabilistically independent of C. And of course this point is essential to practical research into causes. In practice we are never able to identify, let alone gather statistics on, all the multitude of different factors that affect the Es we are interested in. But this doesn't stop us sometimes finding out that some C we can identify is one of the causes of E. For we can be confident of this much whenever we find a positive correlation between C and E that remains even after we hold fixed those other causes of E with which C is probabilistically associated.^{2[2]}

^{2[2]} I do not of course want to suggest that it will be trivial, or easy, or even something we can ever be fully certain about, to identify all other causes X of E that some putative cause C is correlated with. But it is certainly far easier than identifying all causes of E tout court. Moreover, in practice, we can often use background knowledge to attain a fair degree of confidence. Note also how randomized trials exploit the difference between uncorrelated Xs, which we can ignore, and correlated ones, which must be taken into explicit account. The effect of randomizing a 'treatment' C is precisely to push all other Xs into the former category, by forcibly decorrelating them from C. (Cf. Papineau 1989, 1993b.)

The important point in all this is that familiar cases of 'probabilistic causation' are nothing to do with pure quantum mechanical chances. In typical cases where C 'probabilistically causes' E, the known probability of E given C will not correspond to any chance, since C will not constitute a homogeneous reference class for E.

Note what this means for the significance of conditional probabilities. When survey research shows us that $P(E/C)$ is greater than $P(E/-C)$, and that this correlation is non-spurious in the sense that it does not disappear when we condition on further variables, this does not mean that C alone fixes that chance for E. Nor does it even mean that C, in conjunction with whichever other Xs are present in given circumstances, always increases the chance of E by the difference $P(E/C) - P(E/-C)$. For it may be that C interacts with some of these other Xs, making different differences to the chance of E in combination with different Xs, or perhaps even decreasing the chance of E in combination with some special Xs. All the non-spurious $P(E/C) - P(E/-C)$ implies is that C rather than not-C makes that much difference to the chance of E on weighted average over combinations of presence and absence of those other Xs (with weights corresponding to the probability of those combinations).^{3[3]}

3 Pseudo-Indeterminism and Common Causes

Now, these points do not yet constitute an argument for the 'pseudo-indeterministic' thesis that there are always deterministic structures underlying surface probabilities. It is one thing to argue that survey research always involves unconsidered 'error terms' which make further differences to the chances of effects. It is another to hold that, when these 'error terms' are taken into account, the chances of effects are then always zero or one. This would not only require further error terms which make some differences to the chances of effects; in addition, these further differences must leave all chances as zero or one.

Still, I think there is some reason to hold that just such deterministic structures lie behind the causal relationships we are familiar with. This relates to a feature of common causes discussed in the last section. Recall how common causes 'screen off'

^{3[3]} In the view of Ellery Eells, 'Average effect is a sorry excuse for a causal concept' (1991, p. 113). Let me make two points. First, this causal concept doesn't seem so sorry in connection with rational action, since in that context it is exactly what we need. We don't normally know all the details of our situation, and in such cases rationality dictates precisely that we should perform an action C in pursuit of E just to the extent that C non-spuriously increases the probability of E on weighted average over all the situations we might be in. (Beebe and Papineau, 1997; Papineau, forthcoming.) And, second, the average effect concept should not in any case be thought of as a primitive concept of causation, since it is nothing but an average over the different ways C might operate as a single-case cause of E in particular cases. I take C to be such a single-case cause if the chance of E in the actual circumstances is higher than it would have been if C had been absent. (This is not to deny that these single-case counterfactuals might be reducible to generic causal laws and initial conditions; but the generic laws needed here will need to be more fine-grained than statements of C's 'average effect' on E.)

correlations between their joint effects: the joint effects will display an initial unconditional correlation, but conditioning on the presence and absence of the common cause renders them uncorrelated.^{4[4]}

Before I explain how this screening-off phenomenon bears on the issue of pseudo-indeterminism, it will be useful to digress for a few paragraphs, and first consider how screening off illustrates the temporal asymmetry of probabilistic cause-effect relationships. Note how the probabilistic screening-off relation between common causes and joint effects is absent from the 'causally-reversed' set-up where a common effect has two joint causes. We don't generally find, with a common effect (heart failure, say), that different causes (smoking and over-eating, say) are correlated; moreover, when they are, we don't generally find that the correlation disappears when we control for the presence and absence of the common effect.^{5[5]}

When I first started working on the direction of causation, I thought that this asymmetry might provide the key (Papineau, 1985b). Think of the problem as that of fixing the causal arrows between a bunch of variously correlated variables. Just knowing which variables are pairwise correlated clearly won't suffice, since pairwise correlation is symmetrical--if A is correlated with B, then B is correlated with A. But if common causes differ from common effects in respect of screening-off, then perhaps we can do better, and can mark C down as a common cause of joint effects A and B, rather than an effect of A or B, whenever we find a C that screens off a prior correlation between an A and B.

In fact, though, this is too quick. For screening off does not itself ensure that C is a common cause. The screening-off probabilities are also displayed when C is causally intermediate between A and B (thus $A \rightarrow B \rightarrow C$, or $B \rightarrow C \rightarrow A$). So a probabilistic 'fork' (C screens off A from B), doesn't guarantee that C is a common cause. C could also be causally intermediate between A and B. Still, even this gives us something to work with. When we find a probabilistic fork, we can at least rule out C's being a common effect ($A \rightarrow C \leftarrow B$), and be confident that one of the other three possibilities

^{4[4]} This is in fact an oversimplification. In more complex causal structures, for example where A and B have two common causes, neither common cause will screen off the A-B correlation by itself. (Hausman, 1998, p. 209.) A similar point applies to the screening-off property of intermediate causes. I shall skate over this complication in what follows, as it will not matter to the overall argument.

^{5[5]} It should be noted, though, that in particular cases controlling for the presence (but not absence) of a common effect will have precisely the effect of screening off a prior correlation between joint causes. This is because controlling for a common effect serves to induce a negative correlation between its joint causes (Hausman, 1998, p. 83). If the numbers are right, this effect can thus cancel out a pre-existing correlation between joint causes. (Cf. Irzik, 1996, sect 5.) Irzik raises this point because it disproves a claim about the connection between causes and probabilities that I made in earlier work (1993a). However, it does not, so far as I can see, affect the version of reductionism outlined in section 4 below. For some further explanation of why controlling for common effects induces negative correlations between joint causes, see footnote 8 below.

applies. And then, perhaps, by repeatedly applying this inference to different triples from the overall bunch of variables, we might be able to determine a unique ordering among them all.^{6[6]}

In the end, however, the screening off asymmetry turns out to be less central to the reductionist programme that I originally supposed. In the next section I shall borrow from Dan Hausman's work (1998) to lay out the basic requirements for a reduction of causation to probabilities. From the perspective there developed, the important requirement is not so much that common and intermediate causes should screen off unconditional correlations, but rather that there should be probabilistically independent causes for any given effect. This requirement by itself is enough to tell us, for any correlated A and B, whether A causes B, B causes A, or whether they are effects of a common cause. Relative to this basic independence requirement, screening off only plays the relatively minor role of distinguishing direct from indirect causes (and indeed the screening off property of common causes, as opposed to that of intermediate causes, seems to play no important role at all).

Still, there remains an important connection between Hausman's basic independence requirement and the screening-off property. If we conjoin the independence requirement with the hypothesis of pseudo-indeterminism, then we can explain screening-off, in a way I shall outline in a moment, when otherwise the screening-off phenomenon must be taken as primitive. This returns me to the main theme of this section. I want to argue for pseudo-indeterminism (that is, the existence of underlying deterministic structures), on the grounds that we need pseudo-indeterminism to explain the phenomenon of screening off.

I say we can explain the screening-off phenomenon in terms of pseudo-indeterminism and Hausmann's probabilistic independence condition. It is worth being clear exactly what I am aiming to explain here. My idea isn't to explain why certain causally specified structures display probabilistic screening off. From my perspective, this connection holds as a matter of metaphysical necessity. Though the suggestion has yet to be made good, I am assuming that probabilistic screening-off is part of what constitutes certain events as being causally interposed between others. Rather, my target here is to explain why there should be any events that are related by screening off in the first place--that is, why we should find any triples such that A and B are unconditionally correlated, yet the correlation disappears when we control for C.

At bottom, the idea is simple. Let me articulate it in connection with common causes. I shall return to intermediate causes at the end of the section. Suppose that some common cause C has two effects A and B. Suppose further that there are 'error terms' X and Y, such that CX or $Y \leftrightarrow A$. That is, C is a deterministic 'INUS' condition for A: whenever A occurs, either C&X determines it, or some other Y does. Now suppose similarly that C is a deterministic INUS conditions for B: there are further 'error terms' S and T such that $C&S$ or $T \leftrightarrow B$. Then, provided the relevant error terms are probabilistically independent of each other, this guarantees that there will be an unconditional correlation between A and B, which will be screened off by C.

^{6[6]} This general idea goes back to Reichenbach (1956).

Intuitively, if you've got A, then this adds to the probability that C was there to produce it, and so, since C also makes B more likely, this will generate a resulting correlation between B and A. (More carefully, this inference will be valid precisely insofar as the factors which ensure B in the presence of C--namely, S-or-T--are not negatively correlated with A.)

Then, to get the screening-off, note that, among cases where C is present, X-or-Y and S-or-T will be necessary and sufficient for A and B respectively--and among cases where C is absent, Y and T will be similarly necessary and sufficient for A and B respectively. So, as long as these error terms are appropriately probabilistically independent, A and B will also be probabilistically independent given C and not-C. (Cf. Papineau, 1985b.)^{7[7]}

By way of confirmation of this story, note how it will not work 'in reverse' to predict screened-off correlations among deterministic joint causes of common effects. For in these 'backwards' cases it is clear that the background conditions will not satisfy the requisite independence requirements. To return to our earlier example, it is not impossible that there should be some X and Y such that a heart attack (H) plus X, or Y, is coextensive with earlier smoking. (Think of X as signs that the heart attack comes from smoking, and of Y as non-heart-attack traces of smoking.) Similarly, H&S or T might be necessary and sufficient for earlier overeating. But we won't on this account expect to find that overeating is correlated with smoking, precisely, because given a heart attack, overeating can be expected to be negatively correlated with X or Y, and in particular with X, since in cases of overeating it is less likely that there will be signs that the heart attack has been preceded by smoking.^{8[8]}

Now, one natural reaction to this suggested explanation of common causal screening-off would be 'Sure, that's what we should believe if determinism were true. But determinism isn't true.'

I agree that determinism is not true. Quantum mechanics shows that a full specification of current circumstances will often fail to determine what happens next. Still, it is important that my suggested explanation of common causal screening-off does not presuppose universal determinism, according to which every occurrence is fixed by prior circumstances. Indeed, the requirements of my explanation are so weak that it is somewhat misleading to characterise them as 'deterministic' at all. All the explanation requires is that there be further factors which, together with C, fix the

^{7[7]} Note how the 'error terms' will not be probabilistically independent, and so not ensure screening off, in the case where there is another common cause for A and B apart from C (cf. footnote 4). For in these cases some of the non-C-involving causes of A and B will themselves be correlated, due to this further common cause.

^{8[8]} This now helps us to understand why controlling for the presence of common effects should induce negative correlations among joint causes (cf. footnote 4). In the example, overeating becomes a negative indicator of smoking, once the heart attack is taken as given--for once we know that the heart attack came from overeating, we have less reason than before to suspect the victim of smoking.

occurrence of A and B respectively, and that these further factors should be probabilistically independent. It does not require that these further factors, or C for that matter, should themselves be determined by anything. Nor indeed does it require that these further factors be contemporaneous with C: they could equally well be circumstances which emerge chancily in the interval between C and its effects.

Still, this is unlikely to satisfy my objector. Maybe my story doesn't imply universal determinism. But I can still be pressed on why we should suppose even the limited kind of deterministic structure at issue (determination of the effect by the explicit cause plus hidden factors). I earlier made the point that the familiar logic of survey research commits us to the existence of unknown 'error terms'. But, as I admitted at the start of this section, this in itself falls short of the further contention that these 'error terms' will always fix prior chances of zero or one for effects. And as yet I have offered no real evidence for this further contention. Why suppose that all screening-off common causes must be associated with 'error terms' in conjunction with which they determine their effects? After all, if quantum mechanics shows that some earlier events fix chances other than zero or one for later events, then surely I should allow that some screening-off common causes will similarly relate in a purely chancy way to their joint effects? My earlier points about survey research seem to leave it quite open that there should be purely 'quantum mechanical common causes', Cs which fix pure chances for correlated joint effects A and B.

I agree that this possibility is open. At the same time, it is one of the most striking features of quantum mechanics that, when we look for cases of prior quantum mechanical states that fix pure chances for correlated events A and B, we find that such prior quantum states characteristically fail to screen off these correlations, and so fail to display the probabilistic structure constitutive of common causes. (What is more, it can be shown that many such cases involve structures of correlations which cannot be screened-off by any local prior states, even if they are different from those prior states recognized by orthodox quantum mechanics.)

I am here thinking of the well-known 'EPR' correlations. This is not the place to add to the large literature on this subject. My point is simple enough. Once quantum mechanics persuades us of indeterminism, it seems natural enough to suppose that some prior Cs will simultaneously fix pure chances for two events A and B, in such a way that A and B are correlated and their correlation will be screened off by C. But, in fact, when we look into quantum mechanics, we don't normally find pure-chance-fixing prior states with this screening-off characteristic. Moreover, this can't be put down to the temporary failings of current quantum theory, since the relevant correlations often have a collective correlational structure which cannot possibly be screened off by any local prior states.

There isn't of course anything conceptually impossible about some C on its own fixing pure chances for A and B, such that A and B are correlated and C screens off the correlation. But what the EPR cases bring home is that there is nothing conceptually inevitable about this either. That is, it is perfectly possible to have some prior state C which fixes pure chances for A and B, such that A and B are correlated, but where C doesn't screen off this correlation.

Compare this with the deterministic case where C&X or Y determine A, and C&S or T determine B. Given the independence requirements on the 'error terms', this deterministic set-up guarantees that C will screen off the correlation between A and B. In the purely chancy case, by contrast, there is nothing which forces C to screen-off the A-B correlation.

If we look at things from this perspective, it is possible to be less surprised than many people are by the fact that quantum mechanics gives rise to strange correlations that can't possibly be screened off by common causes. Think of it like this. If a prior factor C fixes two later results via some underlying deterministic structure, then, given independence requirements, there is no possibility of the correlation between the later results not being screened off by C. On the other hand, if we don't have this underlying deterministic structure, as in purely chancy quantum mechanical situations, then there is probabilistic room, so to speak, for A and B to become correlated in a way that isn't screened off by prior circumstances. And it turns out that, as soon as nature has this room, it uses it to produce unscreenoffable correlations. What's so surprising about that?

So far I have only defended this pseudo-indeterministic explanation of screening off in connection with common causes. But the same story could be told for intermediate causes ($A \rightarrow C \rightarrow B$). Provided there are 'error terms' together with which such Cs determine their Bs, and provided these are probabilistically independent of backwards 'error terms' which fix whether or not the Cs come from the As, then we will have an explanation for the screening-off property of intermediate causes quite analogous to that given for common causes, and with all the same virtues.

Perhaps it is worth emphasising once more that these 'pseudo-indeterministic' explanations only require a very minimal kind of determinism. What we need are laws of the form $C \& X \vee Y \leftrightarrow A$, where X and Y are independent of the similar factors together with which C fixes other effects and causes. As I said, this doesn't require that X or Y or C are themselves determined. Nor does it require that X or Y be contemporaneous with C. It would be enough, for instance, if there were a plethora of chancy microscopic occurrences temporally between C and A, some of which 'helped' C to fix A, some of which 'hindered' this, and some of which had the power together to fix A even without C. Provided these microscopic events plus C collectively fix whether or not A occurs, and provided they are probabilistically independent of similar factors relevant to the other effects and causes of C, then we will get the probabilistic screening-off structure.

4 Bayesian Nets

So far I have argued for a certain metaphysical picture. Behind the variables involved in normal examples of probabilistic causation are further error terms which also make a difference to the chances of effects. Survey research can ignore these error terms, provided they are probabilistically independent of the causes under explicit study. Moreover, if we assume that these independent error terms suffice to determine effects and causes, then we can explain the screening-off property displayed by common and intermediate causes. Since quantum mechanics suggests that this screening-off property would not be displayed without such deterministic structures,

we thus have reason to accept that such structures underly the apparent indeterminism of familiar probabilistic causes.

In the rest of this paper I want to use this metaphysical picture to respond to some standard objections to the reductionist programme. A first task, however, is to explain how this reduction might work. So far I have alluded to the possibility of reducing causes to probabilities, without actually explaining this.

We think of the problem like this. Suppose we have some set of variables V , together with their joint probability distribution. We know about the correlations between any pair of variables in V conditional on any others. The question is whether this information serves to fix a causal order among those variables. Do the correlations determine which variables should be linked by direct causal arrows?

There has been a great deal of work on this problem over the past decade under the heading of 'bayesian net'^{9[9]} theory, particularly from research groups led by Judea Pearl and Clark Glymour respectively (Sprites, Glymour and Scheines; 1993; Pearl and Verma, 1994). As I pointed out earlier, it is clear that not any correlations among any set of variables will serve to fix a causal direction. (If C screens off an initial correlation between A and B , then this in itself won't decide between: (i) A causes B through C , (ii) B causes A through C , or (iii) A and B are joint effects of C .) Still, the 'bayesian net' research has shown that such causal ambiguities can always be resolved by more complex sets of correlations, perhaps involving further variables. (Cf. Hausman, 1998, pp. 211-4.)

Such inferences from correlations to causes hinge on assumptions about the relationships between correlations and causes. In the literature, these come under various titles, the most familiar of which is the 'Causal Markov condition' of Spirtes, Glymour and Scheines. Here I would like appeal to a codification of relevant assumptions developed by Dan Hausman in his Causal Asymmetries.

As it happens, Hausman himself doesn't aim to reduce causes to correlations, but to a primitive notion of 'nomic connection', for reasons to which I shall turn in later sections. However, those who do seek a reduction to correlations, like myself, will do well to mimic Hausman's elegant reductive strategy, simply substituting 'probabilistic correlation' for 'nomic connection' where necessary. The resulting theory does a great deal to clarify the issues. In particular, if we follow Hausman, we can proceed in stages, first specifying assumptions which allow us to move from correlations to decisions about which variables causes which, but which don't decide whether such causal links are direct or indirect. Further assumptions then allow us to discriminate direct from indirect causes.

^{9[9]} I am unclear about the rationale for this terminology. It is not obviously appropriate, given that the subject has nothing to do with personal probabilities as such, nor with updating them by conditionalization. True, such updating of personal probabilities provides one good theory of statistical inference. But, as explained in footnote 1, questions of statistical inference are best kept separate from questions about inferring causes from lawlike probabilities.

When we divide things up in Hausman's way, it turns out that the basic source of causal direction is not the screening-off asymmetry on which I originally focused, but simply the requirement that effects should have probabilistically independent causes. This latter requirement is all we need at first pass, when we are only aiming to decide what causes what. Probabilistic screening off only matters at the second step, when we need to ascertain whether A causes B directly or through some intermediary.

Hausman shows (modulo my substitution of 'correlations' for 'nomic connections') that if we assume

(I) A and B are correlated if and only if either A causes B, or B causes A, or they have a common cause

and

(II) if A causes B, or A and B have a common cause, then B is correlated with something that is probabilistically independent of A

then it follows that

(III) A causes B if and only if A and B are correlated, and everything correlated with A is correlated with B, and something correlated with B isn't correlated with A.

The two premises involved here will be examined in detail in the following two sections. But I trust they strike readers as having at least some intuitive appeal.

The first premise (I) is simply the idea that any correlation must have a causal explanation, combined, in the other direction, with the thought that causal connections will show up in correlations.

The second premise (II) may be less familiar, but also has some intuitive plausibility. It requires only that whenever some A causes some B, or they have a common cause, there will always be some further influence on B (think 'error terms') which is probabilistically independent of A.

Together these two premises mean that we can tell effects from causes simply by noting that the effects have independent sources of variation. Given an A and B that are correlated, B can't cause A if it co-varies with something which is probabilistically independent of A. To see why this works, note that all the factors correlated with a given cause (any of its causes or effects or symptoms) will be correlated with any further effects it has. So B can't possibly cause A, if it covaries with something which isn't correlated with A.

Once Hausman has fixed arrows of direct-or-indirect causation in this way, then it is fairly straightforward to decide which of these arrows are direct and which indirect. Here we need only assume that, if A is only an indirect cause of B, then its initial correlation with B will be screened off by the conjunction of the other causes of B.

Conversely, if the A-B correlation is not so screened off by B's other causes, then A must be a direct cause of B.^{10[10]}

Hausman's story thus offers an explicit reduction of the relevant causal relationships. Conclusion (III) above specifies a necessary and sufficient condition for direct-or-indirect causation solely in terms of correlations. And, given this, we can then explicitly specify that A directly causes B if and only if it directly-or-indirectly causes it and the A-B correlation isn't screened off by any of B's other direct-or-indirect causes. (Note that the reference to 'direct-or-indirect causes' in this latter specification can be eliminated via (III)).

Now, this reduction is only as good as the assumptions from which it follows. If these assumptions are doubtful, then so is the reduction. In the next section I shall look at assumption (II), which specifies that effects B always have sources of variation that are independent of given causes A. In the following section I shall examine (I), which requires causal connections always to be manifested in correlations, and vice versa.

5 Including the Right Variables

According to (II), whenever A causes B, there is some further X which varies with B but not A. Is this generally true?

If we approach this issue with methodological spectacles on, (II) can seem highly problematic. Suppose that we are conducting some survey, and have chosen to focus on some specific set of variables V. Then, for Hausman's principles to deliver the right answers, B must be represented as being correlated with some A-independent source of variation, for all A and B in V where B is not a cause of A. But there is no reason to suppose that this requirement will automatically be satisfied. Even if reality contains such independent sources of variation, our chosen set V may simply fail to include them. The upshot may be that Hausman's techniques fail to determine a causal order among the variables in V, or, even worse, that they determine a causal order which is different from the one that obtains in reality.

This might seem to undermine the reductive promise of Hausman's strategy. For now it seems that his relationships of probabilistic dependency and independency are only guaranteed to fix the right causal order for certain selections of variables V, namely, those which include independent sources of variation for any effects that they

^{10[10]} Note how this only requires screening-off by intermediate causes, not by common causes. As far as I can see, a Hausman-style reduction does not need screening-off by common causes. The assumption that common causes screen off thus puts extra constraints on cause-probability relations, beyond the minimal constraints needed for probabilities to fix causal order. In particular, it implies that a common cause that does not screen off is not a direct common cause. This implication can be methodologically significant in pointing researchers to unobserved causes.

represent.^{11[11]} However, if we have to use unreduced causal notions in this way to specify the conditions under which probabilities will determine causal structure, then the overall Hausmanian package will clearly fail to show that probabilities alone determine causal structure. (Cf. Irzik, 1996, sect 8.)

However, to argue in this way is to confuse methodology with metaphysics. It is true that in practice survey researchers will always work with limited sets of variables V , and that 'bayesian net' methods will therefore lead them astray if these sets do not satisfy the requirement that all effects are represented as having independent sources of variation. But this is a methodological matter. From a metaphysical point of view, all that matters is that reality should satisfy this requirement. It doesn't matter if certain subsets of variables present a misleading picture of causal structure, as long as reality itself does not. The important metaphysical question is whether God can read the causal facts off from the correlational ones, not whether limited human researchers can do this.^{12[12]}

Critics like Irzik are aware that reductionists will seek to defend their programme by switching the focus from methodology to metaphysics. But they see no reason to suppose that the metaphysical realm will fill the gaps exposed by methodological incompleteness. If the V s used by real researchers are not guaranteed to determine the correct causal order, then what reason is there to suppose that simply switching to larger V s will remedy this failing? Might not even God be stuck with an inadequate V ? (Irizik, 1996, sect 7.)

I take the arguments in the first half of this paper to answer this challenge. They show that there are many probabilistically independent 'error terms' behind the macroscopic

^{11[11]} How does this relate to the requirement, stressed in section 3 above, that a survey should include any common causes of the variables under study? Well, showing that a causally prior C screens off a correlation between some A and B is one source of evidence that A and B must have mutually independent sources of variation and so can't be related as cause and effect. But it is an interesting corollary of 'bayesian net' research that we can know this directly of some correlated A and B , even when we haven't identified any common cause, as a result of explicitly identifying independent sources of variation for both A and B . If A varies with something that doesn't vary with B , then A can't cause B ; and if B varies with something that doesn't vary with A , then B can't cause A either.

^{12[12]} In practice, survey researchers standardly add unobservable independent 'error terms' to their original sets V , and thereby specify probabilistic structures which do fix causal structure. But in general they only know where to put these error terms as a result of already knowing that certain variables do not cause others. This, I take it, is the source of the widespread consensus that 'no causes in, no causes out'. Still, this limitation would not apply to a being who could observe the independent 'error terms' directly. (Actually, even survey researchers aren't as badly off as you might suppose, since they often use temporal order to infer that certain variables can't cause others. Metaphysicians will want an account of causal order that does not presume temporal order--Papineau, 1985b--but there is no reason why practical researchers should hobble themselves in this way.)

variables studied in real survey research. Moreover, it seems that these error terms need to be sufficiently pervasive to restore a kind of determinism, if the screening off phenomena we observe are to be explicable.

Given this, there seems plenty reason to suppose that God's V, as opposed to those used by limited humans, will satisfy Hausman's requirement (II). On my picture, reality is far more complex than it appears to survey researchers. Alongside the macroscopic variables that are salient to human beings, many unobserved influences also enter into the laws relating different events. Any given variable will thus be correlated with a multitude of others. From a God's-eye point of view, there will be no shortage of variables available to display the independence relationships required to fix causal direction.^{13[13]}

6 Causes and Correlations

Let me now turn to Hausman's other assumption:

(I) A and B are correlated if and only if either A causes B, or B causes A, or they have a common cause.

This can be queried in both directions. First, aren't there correlations which are of no causal significance? Second, aren't there causal connections which fail to manifest themselves as correlations?

It is these worries that make Hausman himself stand off from a full-blooded reduction of causal direction to correlations, and settle instead for a reduction to a primitive notion of 'nomic connection'. In Hausman's view, while nomic connections and correlations normally go together, they don't always do so, and when they don't the link between causes and correlations is broken.

I am not convinced these difficulties are insuperable. In the next two subsections I shall consider the first kind of query--are there correlations which lack causal significance? After that, I shall consider the converse question--are there causal connections which fail to manifest themselves as correlations?

6.1 Bread Prices and Venice Water Levels

^{13[13]} Maybe there is plenty of independent variation in this world. But what about simple worlds, such as a world containing nothing but two hard atoms which collide with each other? My answer (and Hausman's, 1998, pp. 67-8) is that such worlds would lack causal direction. This thought might help readers to understand why effects should always have independent sources of variation. At first sight it might seem a lucky freak that each effect should have some such independent source for each of its causes. But it's not so odd, if that is what makes this effect an effect of that cause. The lucky thing is only that the world should display sufficiently complex patterns of probabilistic dependence and independence. No extra design is needed to link up these patterns appropriately with causes and effects--this linkage simply falls out of the metaphysical essence of causal direction.

Aren't there plenty of obvious correlations which signify no causal connection? Isn't there is a good correlation, for instance, between bread prices in Britain and water levels in Venice? (They have both been steadily rising since records began.) (Sober, 1988.)

One ploy here would be to query whether such cases are genuinely lawlike, as opposed to sampling artefacts (cf. footnote 3). But I shall not push this line. There may be something odd about the correlation between bread prices and water levels, but I am prepared to accept that it is a genuinely projectible pattern which we can expect to hold up in future cases.

Instead I would like to discount this correlation on the grounds that its instances are inappropriately related. In the normal case where we take correlations to be of potential causal significance, the correlations are calculated from the paired values displayed when two variables are spatiotemporally co-instantiated. It is specifically because of this that these paired values are candidates for causal relationships. (My school and my exam results; your school and your exam results; and so on.) By contrast, the values of variables in different instances will normally bear no specified spatiotemporal relationship to each other, so will not raise any questions of causal influence. (If my school is correlated with your exam results, we need a more complicated model, which respects the requirement that only spatiotemporally co-instantiated values of variables are candidates for causal relationships. For example, we could take the instances over which the correlation is calculated to be pairs of appropriately related people.)

So I would suggest that principle (I) should be qualified, so as to claim that only correlations of a certain kind must have a causal backing. The kind in question are those correlations whose instances bear no specific spatiotemporal relationships to each other (which then ensures that the values of variables in one instance won't causally influence the values of variables in other instances).

The bread-Venice example doesn't conform to these requirements. The correlation here is calculated by taking points of time as instances. But these points in time aren't picked out haphazardly, so to speak, but form a definite time series. In line with this, the bread price at one time affects the bread price at the subsequent time, and similarly with the Venice water levels. As a result, it cannot be assumed that the bread-Venice correlation reflects some causal connection between the paired values of these variables at a given time. Rather it reflects the fact that temporally earlier bread prices influence later ones, and temporally earlier water levels influence later ones. Bread prices predict Venice water levels, and vice versa, only because the paired values have independently undergone similar causal histories.

From now on, accordingly, I shall take it as read that we are dealing with correlations whose different instances bear no specified spatiotemporal relationships to each other, and so can't have been generated in this way.

Note in this connection how the bread-Venice example will display no co-variation between bread and water levels beyond the correlation that is already implied by each times series taken separately. That is, there isn't any tendency for the two series to

peak or fall in tandem. If we subtract that correlation between bread and water levels that can be attributed to the temporal succession of instances, then we are left with no correlation at all. (If there were such a residual correlation, then that would indeed point to a common cause.)

6.2 EPR Correlations

The other obvious problem for the claim that all correlations have causal significance is quantum mechanics. In particular, EPR cases certainly involve correlations, but it seems unlikely that these reflect causal connections. Apart from anything else, the correlated events can be spacelike separated, which means that any causal influence between them would be in tension with special relativity.

In fact, Hausman's methods deal with this kind of case quite naturally (Hausman, 1998, sects 12.6-7).^{14[14]} The results on the different wings of an EPR experiment both lack independent sources of variation. Anything correlated with one is correlated with the other, and vice versa. So Hausman's principle (II) implies that neither is an effect of the other, nor do they have a common cause.^{15[15]}

Accordingly, Hausman defines a notion of 'mutual dependence': two events are mutually dependent if everything correlated with one is correlated to the other. And then he modifies (I) accordingly:

(I') A and B are correlated if and only if either A causes B, or B causes A, or they have a common cause, or they are mutually dependent.

Since the reductive principle (III) already required that we have cause-effect relationships just in case one end of the correlation has an independent source of variation, this reductive principle can stand as before.

6.3 Failures of Faithfulness

Let me now consider the converse problem facing reductionism, the possibility that there are causal connections which fail to manifest themselves as correlations. In the literature these have come to be known as 'failures of faithfulness'--the correlations we observe are not faithful to the underlying causal structure.

Now, it seems all too possible that there should be such cases. To repeat an example I have used previously, suppose that drinking cola (C) both stimulates people to exercise more (E), but also causes them to put on weight (W). And suppose further

^{14[14]} Given Hausman's non-probabilistic programme, his worry isn't that EPR cases yield non-causal correlations, but that they yield non-causal nomic connections. Still, his proposed solution is equally available to my probabilistic programme.

^{15[15]} Perhaps there is room for another option here. If screening-off by common causes isn't essential to the reduction of causation (cf. footnote 10), then perhaps we should regard the prior quantum state as a non-screening-off common cause of the measurement results.

that exercise E independently has a negative influence on weight increase W, to just the extent required to cancel out the direct positive influence of C, and leave us with an overall zero correlation between cola C and weight increase W.



Here we have a causal connection between C and W, but no corresponding correlation. Indeed if we were to look only at the correlations here, we would get the impression that C and W were both causes of E, but were themselves causally unconnected. (After all, they are themselves uncorrelated, but both correlated with E.)

Now clearly this kind of case is unlikely. It requires two causal influences to cancel out precisely, and so to leave us with no correlation at all. It thus involves a kind of freak, which can perhaps be ignored for practical purposes.^{16[16]} But this freakiness is no help to the reductionist programme. For a reductionist has to say that such failures of faithfulness are not just unlikely, but metaphysically impossible. And this seems just wrong. There seems no principled barrier whatsoever to two causal influences cancelling out exactly and leaving us with a zero correlation, as in the above example. Maybe this would be a freak, but as long as it can happen, reductionism is in trouble.

Both Irzik and Hausman regard this as the Achilles' heel of the reductionist programme (Irzik, 1996, sect 6; Hausman, 1998, sect 10.4). From their point of view, such examples show that causal connections lie deeper in reality than mere correlations. Non-reductionists can allow that causal connections normally manifest themselves in correlations. Perhaps they can even think of them as dispositions to generate correlations when combined in various ways. But they won't equate the causal connections with actually manifested correlations, precisely because they will take it to be entirely possible that certain combinations of underlying causal dispositions will generate the 'wrong' correlations, as in failures of faithfulness.

In earlier work (1993a) I suggested that this difficulty might be avoided by attending to yet further variables. For instance, if in the above example we could find some Z which is correlated with Weight but not with Exercise, then this would argue that

^{16[16]} It would be foolish, however, to ignore this danger too readily when our information about correlations comes from small samples. Remember (footnote 1) that survey research involves two stages: first, inferences from samples to correlations; second, from correlations to causes. Even if a freakish cancelling out of precise causal influences is needed to undermine the second stage, approximate cancelling out of causal influences can easily mess up the first stage, by preventing a genuine causal correlation from registering as statistically significant.

Weight cannot be a cause of Exercise after all, despite initial indications. And perhaps a wider network of variables will suffice to pin down the correct causal structure uniquely.

Hausman has pointed out, however, that this strategy is unlikely to serve the reductionist's purposes. For, even if such a wider network can yield additional fixes on the underlying causal structure, the original failure of faithfulness will not be removed. If the unconditional correlation between Cola and Weight is zero, then it will still be zero after I have examined various further correlations. So whatever principles I might be using to identify the real causal structure, they cannot include the basic reductionist premise (I) which says that absence of correlation means absence of causal connection. This principle is unequivocally falsified by the failure of faithfulness. At best I must be using something like 'Whenever different aspects of a correlational structure imply inconsistent causal conclusions, postulate as few failures of faithfulness as possible to resolve the inconsistencies'. (The idea here would be that if I did continue denying a causal connection between Cola and Weight, in the face of extra correlational evidence, I would need to postulate even more failures of faithfulness to explain away that extra evidence.)

While this principle of charity about fidelity seems a sensible enough methodological maxim, I agree with Hausman that it is an unlikely basis for a metaphysical reduction of causation. So instead I would like to adopt a different strategy, and appeal to the metaphysical picture developed earlier in the paper to answer the challenge.

The idea here would be to go finer rather than broader. Instead of looking at wider frameworks of variables, including more distal causes and effects of such initially troublesome trios as Cola, Exercise and Weight, we could switch the focus to a more microscopic level, and include factors which mediate causally between these variables (cf. Hausman, 1998, p. 221)^{17[17]}. If the arguments in the first half of this paper are correct, there is every reason to suppose there are many such variables, indeed enough to restore pseudo-indeterminism, and moreover that they will satisfy probabilistic independence requirements on causes. The hope, then, would be that at this level there will be no failures of faithfulness, and that the reductive principles derived from Hausman will suffice to fix the correct causal order.

Still, it is not enough that there should simply happen to be no failures of faithfulness at such more microscopic levels. Reductionism requires that there couldn't be any such failures. And it might seem that I still have no argument here. Won't it still make perfectly good metaphysical sense that there should be some fortuitous

^{17[17]} Hausman also makes the point (1998, p. 215) that a pseudo-indeterminist like myself is in danger of having correlations go degenerate if I include all relevant microscopic variables, since the probabilities will then all go to zero or one. But this is too quick: even if determinism means that some (total) antecedents fix zero and one for consequents, less-than-total antecedents can still fix intermediate probabilities. True, this reductionist response does implicitly appeal to the 'naturalness' of certain kinds. But Hausman himself needs natural kinds (pp. 66-8). Moreover, the reductionist programme is likely to need them elsewhere as well, to respond the kind of challenges raised by Arntzenius (1990) and Price (1996).

cancelling of causal influences at the lower levels, thus violating faithfulness once more?

However, there is no reason for the reductionist to accept that such failures of faithfulness can go all the way down. This might make initial conceptual sense, but the reductionist can insist that, since the causal order is metaphysically fixed by probabilistic patterns at microscopic orders, there is simply no metaphysical possibility left of such a microscopic pattern painting a false causal picture.

From my reductionist perspective, to think that there is always a deeper level of causal influences underlying the probabilities is to mistake a difference of levels for a difference in metaphysical kind. It is true that, if we start at a macroscopic level, then there is a deeper level of reality underlying the probabilities at that level, which can discredit those probabilities as causally misleading. But this isn't because that deeper level consists of something non-probabilistic and causally *sui generis*. Rather it is simply that the probabilities at the micro-levels trump those at the macro-levels as a means of fixing causal order.

Compare the response of the beginning student when exposed to the textbook Humean view that causation is nothing but constant conjunction. 'That can't be right, because it leaves out the continuous mechanisms connecting causes and effects.' One can see what the student means, but this thought on its own is no threat to the Humean position. For, as the teacher will point out, the interesting question is whether the more fine-grained causal links which make up the 'mechanisms' are themselves cemented by anything more than constant conjunction.

Similarly in the present context. You may feel intuitively that there must real causal mechanisms behind the probabilities linking causes and effects. But this feeling in itself is no serious threat to the reductionist position, since reductionists can simply respond that the links in such mechanisms are nothing but probabilistic patterns at a more fine-grained level.

If reductionists take this line, they owe some explanation of which 'levels' are to count in fixing causal order. They must dismiss probabilities at 'higher' levels as not themselves being constitutive of causal order, otherwise they will be left with no answer to failures of faithfulness at that level. But then which levels do count? Gesturing at an ordering into 'higher' and 'lower' levels does not really serve, since it fails to tell us which probabilities are constitutive of causal relationships and which are not.

The natural reductionist answer is that only the lowest level counts. Causal relationships at higher levels are fixed by those at the lowest level. Patterns of correlation can thus be misleading about causal structure at any higher level. But at the bottom level there is no metaphysical room for such failures of faithfulness, since there the causal order is simply constituted by the correlational order.

What if there is no lowest level, if there is no limit to how fine we can cut up our mechanisms? Then reductionists can adopt a limiting procedure. Provided there is an ordering of levels into more or less fine-cut, they can say that causal order is fixed once we reach a level where no lower level's correlations overturn that order.

Unsympathetic readers are likely to feel that the reductionist programme is now resting on speculation. To which I readily concede that there is much that is unclear here. In particular, I would like a better understanding of the interface between purely quantum mechanical situations, where causal order goes fuzzy in ways touched on earlier, and the 'pseudo-indeterministic' world, where events become determinate enough to fall into patterns constituting causal order.

At the same time, I see no need to apologise for my metaphysical commitments. As I said at the beginning of this paper, a reduction of causation to probabilities would be a fine thing, if it were possible. For it would show us how there can be causal direction in a world whose fundamental laws are symmetrical in time. Given this, there seems to me no need for every plank in the reductionist programme to be nailed down firmly. If we can develop a cogent metaphysical picture of the sources of causal asymmetry, and if this picture can be supported by general considerations, of the kind offered in the first half of this paper, then I would say we have a good theory of causal asymmetry.

Bibliography

F. Arntzenius 1990 'Physics and Common Causes' Synthese 82

H. Beebe and D. Papineau 1997 'Probability as a Guide to Life' Journal of Philosophy 94

E. Eells 1991 Probabilistic Causality (Cambridge: Cambridge University Press)

D. Hausman 1998 Causal Asymmetries (Cambridge: Cambridge University Press)

G. Irzik 1996 'Can Causes be Reduced to Correlations?' British Journal for the Philosophy of Science 47

D. Papineau 1985a 'Probabilities and Causes' Journal of Philosophy 82

D. Papineau 1985b 'Causal Asymmetry' British Journal for the Philosophy of Science 36

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

D. Papineau 1993a 'Can We Reduce Causal Direction to Probabilities?' in D. Hull, M. Forbes and K. Okruhlik (eds) PSA 1992 vol 2 (East Lansing: Philosophy of Science Association)

D. Papineau 1993b 'The Virtues of Randomization' British Journal for the Philosophy of Science 44

D. Papineau forthcoming 'Causation as a Guide to Life' in D. Papineau The Roots of Reason, OUP.

J. Pearl and T. Verma 1994 'A Theory of Inferred Causation' in D. Prawitz, B. Skyrms and D. Westerståhl (eds) Logic, Methodology and Philosophy of Science IX (Amsterdam: Elsevier)

H. Price 1996 Time's Arrow and Archimedes' Point (Oxford: Oxford University Press)

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

Sober, E. 1988 'The Principle of the Common Cause' in J. Fetzer (ed.) Probability and Causality (Dordrecht: Reidel)

P. Spirtes, C. Glymour and R. Scheines 1993 Causation, Prediction and Search (New York: Springer-Verlag)

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