Brit. J. Phil. Sci. **52** (2001), 331–346

Venetian Sea Levels, British Bread Prices, and the Principle of the Common Cause

Elliott Sober

ABSTRACT

When two causally independent processes each have a quantity that increases monotonically (either deterministically or in probabilistic expectation), the two quantities will be correlated, thus providing a counterexample to Reichenbach's principle of the common cause. Several philosophers have denied this, but I argue that their efforts to save the principle are unsuccessful. Still, one salvage attempt does suggest a weaker principle that avoids the initial counterexample. However, even this weakened principle is mistaken, as can be seen by exploring the concepts of *homology* and *homoplasy* used in evolutionary biology. I argue that the kernel of truth in the principle of the common cause is to be found by separating metaphysical and epistemological issues; as far as the epistemology is concerned, the Likelihood Principle is central.

1 The principle and the counterexample

Reichenbach's ([1956]) *principle of the common cause* reads as follows:

(P) If events X and Y are correlated, then either X caused Y, Y caused X, or X and Y are joint effects of a common cause (one that renders X and Y conditionally probabilistically independent).[1]

Some time ago I described a counterexample to this principle:

[1] One needs to assume that X and Y are logically independent; otherwise, problems arise. See Arntzenius ([1993]) for discussion.

> Consider the fact that the sea level in Venice and the cost of bread in Britain have both been on the rise in the past two centuries. Both, let us suppose, have monotonically increased. Imagine that we put this data in the form of a chronological list; for each date, we list the Venetian sea level and the going price of British bread. Because both quantities have increased steadily with time, it is true that higher than average sea levels tend to be associated with higher than average bread prices. The two quantities are very strongly positively correlated (Sober [1987], pp. 161–2).

My point was that:

> we do not feel driven to explain this correlation by postulating a common cause. Rather, we regard Venetian sea levels and British bread prices as both increasing for somewhat isolated endogenous reasons. Local conditions in Venice have increased the sea level and rather different local conditions in Britain have driven up the cost of bread. Here, postulating a common cause is simply not very plausible, given the rest of what we believe (*Ibid.*, pp. 161–2).

This counterexample, if it is a genuine counterexample, should not be dismissed as an idle curiosity. For one thing, principle (P) is a consequence of what has come to be called *the causal Markov condition*:

> For every set of causal variables V that contains X and Y and that is causally sufficient for X and Y, if X does not cause Y, then Pr[X| Parents(X) & Y] = Pr[X|Parents(X)].[2]

Here Parents(X) are all the direct causes of X that are members of V, and causal sufficiency means that V contains all common causes of X and Y. This condition is presupposed by the causal modeling frameworks of Spirtes, Glymour, and Sheines ([1993]) and of Pearl ([2000]). If (P) is wrong, there is work to do in tracing out the consequences of this fact.

2 Two attempts to defuse the counterexample

Meek and Glymour ([1994], p. 1006) responded to the Venice/Britain example by pointing out that the problem it exemplifies was discussed by G. Udny Yule ([1926]),

> who attributed such statistical dependencies either to (1) an unobserved common cause or to (2) an unrepresentative sample or to (3) mixing populations with different causal structures and different probability distributions [numbers inserted by me].

[2] Here is Hausman and Woodward's ([1999]), p. 524) argument, slightly modified, for why the causal Markov condition entails (P): Suppose that X and Y do not have a common cause. Then the two-member set V = {X,Y} is causally sufficient for X,Y. Since Parents(X) is empty in V, the causal Markov condition entails that if X does not cause Y, then Pr(X|Y) = Pr(X). If so, the causal Markov condition also entails that if X does not cause Y, and Y does not cause X, and X and Y do not have a common cause, then Pr(X|Y) = Pr(X), which is the contrapositive of (P).

Meek and Glymour think that these replies suffice to show that the Venice/Britain story is no counterexample to principle (P); Spirtes, Glymour, and Sheines ([1993], pp. 57–64) take the same line. I disagree. I claim that option (1) is not true in the example at hand—suppose I am right that there is no common cause of Venetian sea levels and British bread prices, at least none that would explain the observed correlation by rendering the joint effects conditionally independent of each other. By the way, the mere conjunction of Venetian causes and British causes cannot constitute a *common* cause of both effects, on pain of trivializing the principle; to call this conjunction a common cause would mean that there are no separate cause explanations with which a common cause explanation competes (Sober [1987], [1988]). As for possibility (2), I claim that the sample is *not* unrepresentative. Go back and draw more and more measurements from the two centuries in question; the association of values will remain in place. The pattern in the data is not due to sampling error. Venetian sea levels and British bread prices are correlated during this period because of facts about the underlying probabilities; the correlation is not a superficial and evanescent property of the sample one has at hand.[3]

Moving finally to option (3), I agree that this is the right diagnosis, if 'mixing populations with different causal structures' means that the causal processes driving British bread prices have nothing to do with the causal processes behind Venetian sea levels. However, (3) does not save principle (P); quite the opposite—it shows that the principle is mistaken. If we only knew the true causal relationships, we would not need principles like (P) to tell us how to infer causal relationships from probabilities.[4]

Forster ([1988]), Papineau ([1992]), and Hausman and Woodward ([1999]) responded to the Venice/Britain example in a different way. They point out that a strict monotonic increase in the two quantities would mean that *changes* in bread prices and *changes* in sea levels are not correlated. Suppose British bread prices go up in 100% of the relevant years and that the same is true of Venetian sea levels. Since every year marks an increase in both, there

[3] When we ask whether a correlation 'really exists,' we are imagining what would happen if we sampled repeatedly from 'the same' process. With respect to Venice and Britain, I take this question as follows—given the physical processes and initial conditions actually in place, what would happen if we drew repeated data sets? But it is possible to regard the physical processes and initial conditions that actually obtained as constituting just one 'realization' of a process that could have had many other realizations; under this construal, one can reach the verdict that there is no real correlation here, only an accidental association in the data. To do this, one need only regard the initial conditions and processes as themselves subject to variation of a certain sort. There is no reason to think that the latter way of posing the question is the one that is really correct. The result of this insistence would be to conclude that many perfectly real correlations are unreal, including Reichenbach's ([1956]) actors example, which I'll discuss later.

[4] The injunction against 'mixing populations with different causal structures and different probability distributions' is also supposed to rule out mistakes that arise from Simpson's paradox (Spirtes *et al.* [1993]). In this context as well, knowledge of the causal facts cannot be invoked to save (P), on pain of depriving the principle of its epistemic force.

is no correlation between change in one and change in the other. Assuming that the observed frequency of a change reflects the true underlying probability of the change, we have the following equality:

(E) Pr{[Venetian sea levels in year (i + 1) > Venetian sea levels in year i] & [British bread prices in year (i + 1) > British bread prices in year i]} = Pr[Venetian sea levels in year (i + 1) > Venetian sea levels in year i]Pr[British bread prices in year (i + 1) > British bread prices in year i].

This equality is true because the three probabilities mentioned all have a value of 1.0.

Although this point is correct about *changes* in the two quantities, the fact remains that the *absolute values* of bread prices and sea levels are correlated. As I claimed initially, higher than average bread prices *are* correlated with higher than average sea levels:

> Pr[Higher than average sea level in year i & higher than average bread price in year i] > Pr[Higher than average sea level in year i] Pr [Higher than average bread price in year i].

Imagine that ½ the years between 1800 and 1999 exhibit higher than average bread prices and that ½ the years exhibit higher than average sea levels. The frequency of years in which both these events occur is substantially greater than ¼. Thus, the Venice/Britain example *is* a counterexample to the principle as stated.

As an aid to understanding how absolute values can be correlated without changes being correlated, consider the following hypothetical data set, which I concocted for each of eight successive years. In each year, the average Venetian sea level (in some unit) and the average British bread price (in some unit) are specified:

year	Venetian sea level	British bread price
8	31	20
7	30	19
6	29	15
5	28	14
4	25	10
3	24	6
2	23	5
1	22	4

Since there are 8 years, there are 7 transitions from one year to the next. In both Venice and Britain, increases took place in all 7 transitions. If these frequencies are indicative of the underlying probabilities, we may conclude

that *increases are uncorrelated*, since $(7/7) = (7/7)(7/7)$. However, notice that half the years (years 5–8) have above average sea levels and half the years (again, years 5–8) have above average bread prices. The frequency of years in which sea levels and bread prices are *both* above average (years 5–8) is ½. If these frequencies are indicative of the underlying probabilities, the *absolute values are correlated*, since $½ > (½)(½)$.

It is worth noticing that the monotonic increases don't need to be strict for absolute values to be correlated while increases are uncorrelated. Suppose that Venice and Britain each are subject to their own endogenous *stochastic* processes. Let the sea level in Venice change according to its own set of probabilistic rules and let the same be true of the bread price in Britain. For example, suppose that Venetian sea levels have a probability v of going up by 1 centimeter each year and a probability $(1-v)$ of going down by 1 centimeter. Similarly, suppose that British bread prices have a probability b of going up by 1 penny each year and a probability $(1-b)$ of going down by 1 penny. If increase is more probable than decline in each case (i.e. $v,b > ½$), then bread prices and sea levels will probably show gradual, but not exceptionless, upward trends. There probably will be years in which declines occur, but on the whole the two trends will be up. In this case, the *absolute values* of bread prices and sea levels will be correlated, but *changes* in one will not be correlated with *changes* in the other. The equality (E) will be true, but not because $(1.0) = (1.0)(1.0)$.

3 A weakened principle and a new counterexample

There is an obvious reply to my point about Venice and Britain. Although the example does refute principle (P), it does not touch the following principle:

(P*) If events X and Y are correlated and so are changes in X and changes in Y, then either X caused Y, Y caused X, or X and Y are joint effects of a common cause.

Is this weaker principle correct? I don't think so, but Venice and Britain, as I understand them, do not show why. Let us turn, then, to some ideas from evolutionary biology.

When two species share a trait that they inherited unmodified from a common ancestor, the similarity is called a *homology*. When the trait evolved independently in the two lineages, the resulting similarity is called a *homoplasy*. This difference is illustrated in Figure 1. Robins and sparrows both have wings, and this is a homology; robins and bats both have wings, but this is a homoplasy. As Figure 1 suggests, the claim that wings in robins and sparrows are homologous means that their wings trace back to a common origination event; the claim that wings in robins and bats are

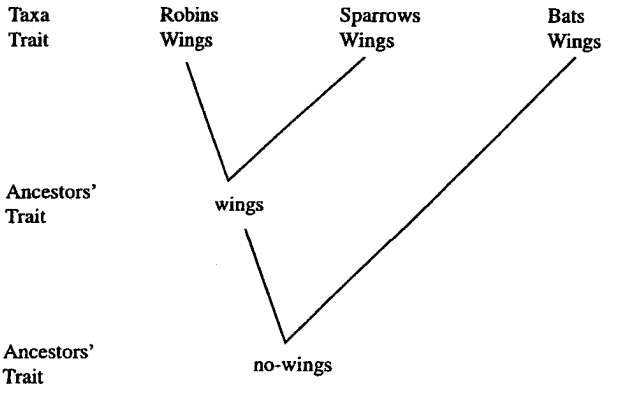

Fig. 1.

homoplasious means that they trace back to separate origination events. Notice that the question is not whether two *taxa* have a common *ancestor*. Robins and sparrows have a common ancestor, but so do robins and bats.

Homologies and homoplasies can involve genetic, morphological, physiological, or behavioral characteristics. Of special interest to us here are traits that involve *ontogenetic sequences*. Organisms such as ourselves go through a sequence of stages as they develop from embryo to neonate to adolescent to adult and from there into senescence. These sequences are often uniform within species and also between them. There are two features of this sort of uniformity that I want to consider. First, the *order* of the stages can be the same, or nearly the same, across a range of organisms; each may develop from stage A to B to C and thence to D. But, in addition, the *timing* of these stages may be similar across organisms. Organisms may stay in stage A for around 9 months, in stage B for the next 12 years, in stage C for the next 7 years, and so on.

If we look at the development of two human beings, we often will see a strong correlation between the stages they occupy and also between their timing of those stages. The natural explanation of this correlation is that it is a homology. The two human beings resemble each other because they inherited the same 'developmental program' from a common ancestor. But now suppose we find similar developmental sequences in two organisms that belong to different species. It is possible that this, too, is a homology, but must it be? The answer is *no*; it is perfectly possible that similar developmental sequences evolved independently in the two lineages.

To make this suggestion more concrete, let's look at a single and relatively simple developmental feature of organisms—their height. Many organisms start small, then gradually increase in height, then go through a growth spurt

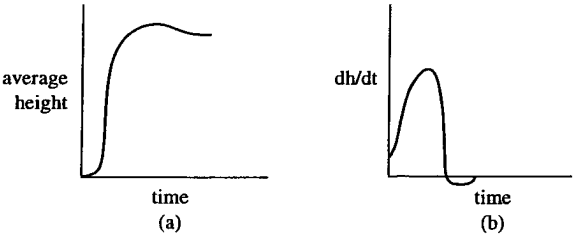

Fig. 2.

in adolescence, and then level off in adulthood. In old age, they shrink a bit. This trajectory of height is depicted in Figure 2a; the rate of change in height is depicted in 2b. Let's think of each line in these two figures as a species average. Individual organisms in the species can have height trajectories that deviate from the species average. But for many (indeed, most) pairs of conspecifics, their heights will be correlated through time and so will their rates of change in height. For two organisms to exhibit this twofold correlation, they need not have the same birthdays and longevities. Just as in the Venice/Britain example, we can list the heights that each achieved at different ages and then show that the heights are correlated; the same is true of the two individuals' rates of change in height.

If these two similarities are homologies, principle (P*) has nothing to worry about. But suppose the two organisms come from different species, whose ancestors independently evolved the developmental sequences depicted in Figure 2. If 'common cause' means common ancestry and 'separate cause' means separate ancestry, then (P*) will be false; the heights attained at different times by the one organism and the heights attained at different times by the other will be correlated (like sea levels in Venice and bread prices in Britain), and changes in height in one organism will be correlated with changes in height in the other (unlike the Venice/Britain example). If the similarities are homoplasies, the explanation of these correlations is in terms of separate causes.

Thus far, I have argued that there are possible exceptions to Principle (P*); if similar developmental sequences evolve independently in two lineages, then the organisms in one and the organisms in the other will exhibit correlations that lead principle (P*) to conclude, erroneously, that this homoplasy is a homology. But how seriously should this possible violation of (P*) be taken? Perhaps homoplasies that involve developmental patterns like the one depicted in Figure 2 are like brains-in-vats and evil demons in traditional epistemology—they constitute epistemological nightmares that we are confident do not afflict us in our waking lives. However, I now want to argue that homoplasies that refute principle (P*) are not just *conceivable*;

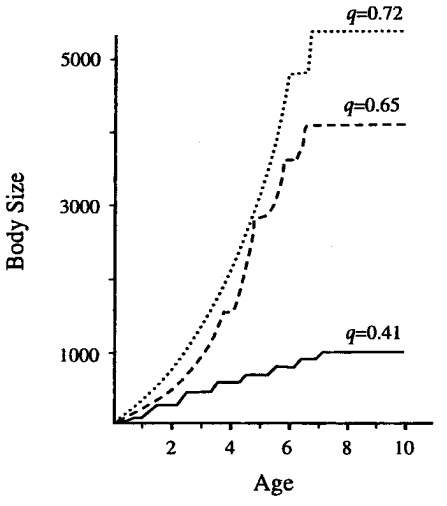

Fig. 3.

there is a body of theory concerning *the evolution of life history strategies* that shows how natural selection can be expected to produce precisely this outcome. Such counterexamples are not just *conceivable*; they are *predictable*.

As an example, consider a plant in a stable population that lives in a seasonal environment. It reproduces at most once each year, but to do so it has to stop growing. Assume further that adult mortality is independent of size, and that the organism has a certain maximum lifetime, say 10 years. Given these and a few other assumptions, Kozlowski and Uchmanski ([1987]) constructed a model from which one can deduce the optimal allocation of energy to growth and reproduction as a function of the organism's age and its probability (q) of surviving to the next year. Optimal body size increases with age, and the organism should devote more energy to reproduction, and less to growth, as the years pass. An organism with a greater risk of mortality should start reproducing sooner, and should shift more quickly from growth to reproduction, than an organism with a smaller risk. Figure 3 (taken from Stearns [1992], p. 165) illustrates their results for three possible values of q.

Any lineage satisfying the model's assumptions will experience a selection pressure that pushes organisms to conform to the type of growth curve depicted in Figure 3. If this selection pressure is strong, and the nonselective forces that could 'get in the way' (Sober [1993], Ch. 5) are relatively weak, we expect the organisms in the population to closely approximate their optimal growth curves. This means that if two lineages have the same value for q and the same maximum life span, selection will cause them to evolve the same growth curves. And even if their values for q and their maximum lifespans are

different, their sizes will be correlated and so will their rates of change in size. In the circumstances described, selection will produce homoplasies that lead principle (P*) into error.

One might try to rescue principle (P*) by saying that organisms that exhibit a homoplasy do so because of *similar* causal influences. Perhaps natural selection acting on one ancestral population led human beings to obey the pattern in Figure 2 and the same sort of selection process acting on another ancestral population led anteaters to do the same. That is, a common *type* of cause impinged on both lineages. This may be so, but it need not be. It is perfectly possible for the same trait to develop independently in different lineages for very different evolutionary reasons. For example, one lineage might evolve a pattern of coloration to attract mates, while another might evolve that pattern for purposes of camouflage. But even if the same *type* of selection pressure occurred in each lineage, that would not rescue the principle of the common cause, which is about common *token* causes. Let me explain this point by considering one of the nice examples that Reichenbach ([1956]) used to describe what his principle means. This is the case of the two actors in a traveling troupe who have correlated sick days. The correlation exists because the actors take their meals together; when one gets poisoned by the spoiled food in a greasy spoon, the other usually is there too, eating the same food. On each day, the two actors are usually in the same token restaurant, not merely in two restaurants of the same type. Of course, this token object has properties—types—that stand in interesting probabilistic and causal relationships to the types exhibited by the two actors. The restaurants have probabilities of serving tainted food, and the actors have probabilities of getting sick, conditional on their meal's being tainted. Actors and restaurants are tokens, while *serving tainted food* and *having gastro-intestinal distress* are types. A common cause explanation postulates a common *token* object; the common cause consists of that object's having a certain suite of properties (Sober [1987], [1988]). This is why homoplasies induced by numerically distinct but qualitatively similar selection processes do not have common cause explanations.

4 Further discussion

It happens to be a feature of Reichenbach's example that the correlation one is led to think about concerns whether the two actors get sick on the same days. However, the simultaneity of the two events is not essential. Suppose that when the two actors go into a restaurant, X eats a meal there while Y eats nothing, but orders a meal for take-out, which she consumes the next day. The upshot is that there will be a correlation between X's being sick on day i and Y's being sick on day (i + 1). Reichenbach's principle (P) applies to

this case as well. In similar fashion, we can apply the principle to two organisms that follow the same developmental program, but are not born on the same day. And even when the two organisms have different life spans, we can map stages in the first organism's life onto stages in the second's. One natural mapping to consider preserves temporal order; if organism X reaches stage one before it reaches stage two, then organism Y reaches the image of stage one before it reaches the image of stage two.

The acting troupe example illustrates Reichenbach's idea that a common cause explanation explains an observed *dependency* between X and Y by introducing a variable that renders X and Y conditionally *independent*. The following equality helps explain the following inequality:

> Pr(X is sick on day i & Y is sick on day i|X and Y eat in the same restaurants on day i)
> = Pr(X is sick on day i|X and Y eat in the same restaurants on day i)Pr(Y is sick on day i|X and Y eat in the same restaurants on day i).
> Pr(X is sick on day i & Y is sick on day i) > Pr(X is sick on day i)Pr(Y is sick on day i).

Given this, one cannot criticize the Venice/Britain counterexample to principle (P) by pointing out that Venice and Britain are independent of each other once one conditionalizes on the separate causes affecting each. They *are* conditionally independent, but this does not mean that sea levels and bread prices are 'really' uncorrelated. Their conditional independence is real, but so is their unconditional correlation. Venice and Britain are correlated in the same sense that the two actors are.[5]

One objection that is sometimes raised to the Venice/Britain counterexample is that the properties of Venice and Britain that I describe are not *intrinsic*. I said that higher than average sea levels are associated with higher than average bread prices, but whether a year has a higher than average sea level or bread price depends on what happens in other years, not just on what goes on in the year itself. Does this mean that principle (P) might be correct when it is restricted to intrinsic, nonrelational, properties?[6] Apparently, the principle yields sensible results when applied to the example of the acting troupe, and here the correlated properties seem to be intrinsic—whether an actor is sick on a given day does not depend logically on whether he or she is sick on other days. One thing that should warn us off this proposal is the fact that it is possible for two species to have correlated intrinsic properties where the shared properties are homoplasies. But more fundamentally, it needs to be seen that the correlated properties of Venice and Britain *are* intrinsic, however much my talk of 'higher than average' sea levels and bread prices

[5] For discussion of the meaning of 'correlation' in principle (P), see Berkovitz [forthcoming].
[6] I am grateful to Joseph Berkovitz for raising this question.

might suggest otherwise. To see why, it is helpful to focus first on the acting example. Here we have a dichotomous character—each actor is either sick or not sick on a given day—and we express the idea of correlation or nonindependence by describing the probabilities of dichotomous propositions. This idea can be expressed equivalently by talking about a characteristic function that takes the value 1 when an actor is sick on a given day and 0 when he or she is not. The expected value of this function is just the expected frequency of sick days—i.e. the underlying probability of being sick on a randomly selected day. If the two actors' sick days are correlated, this predicts that days on which one actor has a higher than average value (i.e. a value of 1) will be associated with days on which the other actor has a higher than average value. To say that an actor has a higher than average (= expected) value on a given day is just a roundabout way of saying that he or she is sick; it does not entail that the correlated characters are not intrinsic. The only difference between the traits involved in the actor example and those considered in the case of Venice and Britain is that the first are dichotomous while the second are quantitative. Sick days are correlated in the first example and sea levels and bread prices are correlated in the second. In both cases, the correlated properties are intrinsic, though what it means for them to be correlated involves the relationship of what happens at one time with what happens at others.

Another suggestion that might be contemplated as a device for saving Reichenbach's principle is to stipulate that it should not be applied to time series data at all. The problem with this restriction is that it prevents the principle from capturing the fact that correlated ontogenetic sequences are often plausibly explained by the hypothesis of common ancestry. The fact that two human beings have similar ontogenies *is* evidence of common ancestry, and an adequate principle of the common cause should capture this fact. Indeed, if we want the principle to apply to Reichenbach's actor example, we should reject this suggestion, since that example involves a time series; on each day, we record whether each actor is sick.

Yet another proposal for saving Reichenbach's principle is that it should be limited to correlated events that are spatio-temporally contiguous in some special way that Venice and Britain do not exemplify. It is hard to know how to state this restriction exactly; after all, the two actors are not belly to belly as they wend their ways from restaurant to restaurant. And the points so far developed about homologies and homoplasies further suggest that this line of thinking will not bear fruit. If two species exhibit homologies, the spatial and temporal distance between them may be precisely the same as two other species that exhibit homoplasies. Of course, when correlated events *do* trace back to a common cause, there is a spatio-temporal contiguity (in fact, an identity); the actors' sick days trace back to their being in the same token

restaurants at the same times, and the wings of sparrows and robins trace back to the token wings in a single ancestral population. But this is something we know only after we have accepted the common cause explanation; spatio-temporal contiguity is not a feature of the data that helps us choose a plausible explanation.

Should we replace (P*) with a still weaker principle (P**) in which we assert that X causes Y or Y causes X or X and Y are joint effects of a common cause if the absolute values, the first derivatives, and the second derivatives of X and Y all are correlated? I don't think so. Homoplasies of even this intricate sort are not impossible. Correlations in a time series involving as many derivatives as you please can be due to separate causal processes.

5 The rational kernel

Although correlations need not have explanations of the sort described in principle (P), such explanations often are plausible. When is that so? Part of the answer is simple—*when they compete with separate cause explanations that predict that correlations do not exist.* In the acting troupe example, the common cause hypothesis competes with a separate cause explanation. The former predicts a correlation of sick days, while the latter predicts that the sick days should be uncorrelated. If so, the observed correlation favors the former hypothesis over the latter. In the Venice/Britain example, the common cause explanation predicts a correlation of sea levels and bread prices, but so does the obvious separate cause explanation. If so, the observed correlation does not discriminate between the two hypotheses. However, if the common cause explanation predicts that *changes* in bread price will be correlated with *changes* in sea level, while the separate cause explanation says that these changes will occur independently, the data can provide a basis for choosing between the two hypotheses. And when we turn to the developmental sequences discussed in biology, the choice between common cause and separate cause explanations depends on finding predictions about which the two models disagree. If both predict a correlation of heights, the observed correlation is useless. And if both predict a correlation of changes in height, then that observed correlation is useless as well. In each of these examples, the relevant epistemological maxim is the *Likelihood Principle* (see Edwards [1972]; Royall [1997]):

> Observation O favors hypothesis H1 over hypothesis H2 precisely when $\Pr(O|H1) > \Pr(O|H2)$.[7]

[7] For discussion of why likelihood should be thought to measure evidential support, see Forster and Sober ([forthcoming]).

When X and Y are observed to be associated in the data, a common cause explanation will often say that the association was to be expected, while a separate cause explanation will often say that the association is surprising. If so, the common cause explanation is more likely.

The principle of the common cause is supposed to provide epistemological advice, but the principle is constitutionally unsuited to that task. The reason is that the principle has the wrong *form*. To explain what I mean here, I will use the term 'association' to describe a feature of the actual frequencies exhibited by one's data and 'correlation' to describe a feature of the underlying probabilities. Although Reichenbach discussed inferring causal relationships from observed 'correlations,' it is better to think of the process as having two steps, not just one. Correlations are inferred from associations; then, on the basis of these inferred correlations, one tries to infer a causal model. The first formal defect in principle (P) is that one cannot decide whether a common cause explanation is well supported by a correlation unless one knows what the competing hypotheses are (Sober [1999]). This contrastive element is absent in the P-principles I have considered, but comparison is of the essence if we want to know which model is better supported by the data.[8] If a common cause model deductively entails that a correlation exists, while a separate cause model deductively entails that no correlation exists, then, if you think that the correlation *does* exist, this conclusively decides the matter in favor of the common cause hypothesis. However, this veneer of conclusiveness disappears when we consider the bearing of an observed *association* on the problem of model selection. The common cause model does not entail that the observed association must exist, and the separate cause model does not entail that the observed association cannot exist. Rather, the two models confer different probabilities on the observed association. Yet, principles (P) and (P*) both take the form of nonprobabilistic if/then statements.

Once these two changes are introduced into the principle of the common cause—explicit mention is made of competing hypotheses and it is recognized that hypotheses do not entail observations, but confer probabilities on them—we can identify the kernel of truth that the principle contains. This rational kernel is the idea that common cause explanations often disagree with separate cause explanations over whether associations should exist. When they *do* disagree in this way, the observed association favors, but does not conclusively prove, the common cause hypothesis.

[8] Berkovitz ([forthcoming]) points out that one of Reichenbach's ([1956]) formulations of his principle says that 'improbable correlations' are to be explained in one of the three ways described by principle (P). This formulation may be interpreted charitably as a likelihood claim—that associations that have low probability on the hypothesis of separate causes, but high probability on the hypothesis of common cause, favor the latter hypothesis.

6 The metaphysical converse

The Likelihood Principle says nothing about the concept of cause, but the principle of the common cause gives advice about inferring causal models from correlations (or better, from observed associations). This suggests that there must be more to the principle of the common cause than the Likelihood Principle; what is needed is some principle that describes how cause and correlation are connected. If (P) is wrong, something else must be found to do the job.

There *is* something more. Just as (P) may be thought of as an epistemological principle, its converse may be interpreted as a 'metaphysical' claim. The converse principle says that if X causes Y, or X and Y are effects of a common cause, then X and Y will be correlated.[9] This metaphysical principle also contains a kernel of truth, but isolating that kernel likewise requires that some qualifications be introduced. If there is just one causal pathway connecting X to Y, then X and Y will be correlated. But if there are multiple pathways, X can be positive for Y along one pathway and negative for Y along the others, with the result that X and Y are uncorrelated overall. Similarly, if C is a common cause of X and Y (and X doesn't cause Y nor does Y cause X), then X and Y will be correlated, unless there exists a second common influence $D \neq C$ on X and Y that nullifies the correlation that C by itself would induce.

These two metaphysical principles describe what various causal models predict about probabilistic relations; these predictions can then be checked by looking for patterns of association in the data. They have nothing to fear from Venice and Britain, nor from evolutionary homoplasies. As far as I know, they are not impugned by quantum mechanical considerations. The reason I mention them, however, is not so much to claim that they are exactly correct, but to point to the need for principles that describe the probabilistic predictions that various causal models entail. Given a battery of such principles, there is no need to search for a suitably restricted version of the principle of the common cause. The Likelihood Principle, and other relevant epistemic considerations,[10] which have nothing specifically to do with causality, can then be invoked. The *metaphysics* of causation, of course, is specifically about the nature of causality; but given an appropriate metaphysics, the *epistemology* of causation need have nothing specifically to do with the concept of cause.

[9] This converse principle is a consequence of what Spirtes *et al.* ([1993]) term 'faithfulness.'

[10] The Likelihood Principle is not the whole story. For example, when the hypotheses of common and separate causes have objective prior probabilities, these need to be taken into account. And when causal models contain different numbers of adjustable parameters, the Akaike information criterion and related criteria of model selection become relevant. See Forster and Sober ([1994], forthcoming).

Acknowledgements

I am grateful to Joseph Berkovitz, Nancy Cartwright, Ellery Eells, Malcolm Forster, Clark Glymour, Daniel Hausman, Christopher Hitchcock, David Papineau, James Woodward, and to the referees of this journal for useful suggestions. I also benefited from discussing this paper at the London School of Economics and Political Science, at Oxford University, and at University of Wisconsin-Madison.

Department of Philosophy
University of Wisconsin
Madison, WI 53706
USA
ersober@facstaff.wisc.edu
Department of Philosophy, Logic, and Scientific Method
London School of Economics and Political Science
Houghton Street
London WC2A 2AE
UK

References

Arntzenius, F. [1993]: 'The Common Cause Principle', in D. Hull and K. Okruhlik (*eds*), *PSA 1992*, Vol. 2, East Lansing, MI: Philosophy of Science Association, pp. 227–37.

Berkovitz, J. [forthcoming]: 'The Many Principles of the Common Cause', *Reports on Philosophy*.

Edwards, A. [1972]: *Likelihood*, Cambridge: Cambridge University Press.

Forster, M. [1988]: 'Sober's Principle of Common Cause and the Problem of Comparing Incomplete Hypotheses', *British Journal for the Philosophy of Science*, **55**, pp. 538–59.

Forster, M. and Sober, E. [1994]: 'How to Tell When Simpler, More Unified, or Less *Ad Hoc* Theories Will Provide More Accurate Predictions', *British Journal for the Philosophy of Science*, **45**, pp. 1–36.

Forster, M. and Sober, E. [forthcoming]: 'Why Likelihood?' in M. Taper and S. Lee (*eds*), *The Nature of Scientific Evidence*, Chicago, IL: University of Chicago Press. Also available at the following URL: http://philosophy.wisc.edu/forster.

Hausman, D. and Woodward, J. [1999]: 'Independence, Invariance, and the Causal Markov Condition', *British Journal for the Philosophy of Science*, **50**, pp. 521–83.

Kozlowski, J. and Uchmanski, J. [1987]: 'Optimal Individual Growth and Reproduction in Perennial Species with Indeterminate Growth', *Evolutionary Ecology*, **1**, pp. 214–30.

Meek, C. and Glymour, C. [1994]: 'Conditioning and Intervening', *British Journal for the Philosophy of Science*, **45**, pp. 1001–21.

Papineau, D. [1992]: 'Can We Reduce Causal Direction to Probabilities?' in D. Hull and K. Okruhlik (*eds*), *PSA 1992*, Vol. 2, East Lansing, MI: Philosophy of Science Association, pp. 238–52.

Pearl, J. [2000]: *Causality: Models, Reasoning, and Inference*, Cambridge: Cambridge University Press.

Reichenbach, H. [1956]: *The Direction of Time*, Berkeley: University of California Press.

Royall, R. [1997]: *Statistical Evidence: a Likelihood Paradigm*, Boca Raton, FL: Chapman & Hall/CRC.

Sober, E. [1987]: 'The Principle of the Common Cause', in J. Fetzer (*ed.*), *Probability and Causation: Essays in Honor of Wesley Salmon*, Dordrecht: Reidel. pp. 211–28. Reprinted in *From a Biological Point of View*, Cambridge: Cambridge University Press, 1994, pp. 158–74. [page reference to the latter]

Sober, E. [1988]: *Reconstructing the Past: Parsimony, Evolution, and Inference*, Cambridge, MA: MIT Press.

Sober, E. [1993]: *Philosophy of Biology*, Boulder, CO: Westview Press, 2nd edn 1999.

Sober, E. [1999]: 'Testability', *Proceedings and Addresses of the American Philosophical Association*, **73**, pp. 47–76. Also available at the following URL: http://philosophy.wisc.edu/sober.

Spirtes, P., Glymour, C., and Sheines, R. [1993]: *Causation, Prediction, and Search*, New York: Springer Verlag; 2nd rev. edn, Cambridge, MA: MIT Press, 2000.

Stearns, [1992]: *The Evolution of Life History*, Oxford: Oxford University Press.

Yule, G. U. [1926]: 'Why Do We Sometimes Get Nonsensical Relations Between Time Series? A Study of Sampling and the Nature of Time Series', *Journal of the Royal Statistical Society*, **89**, pp. 1–64.