

Two Concepts of Cause¹

Elliott Sober

University of Wisconsin, Madison

1. Introduction

Hume contributed two ideas to our understanding of causality. First, he held that when one event causes another, the physical ontology of this situation contains two items, not three. There is the causing event and its resulting effect; but there is not, in addition, a third physical entity consisting of their causal relation. Causality, Hume believed, is a relation like that of one individual's being taller than another. Sam and Aaron are terms of this relation. Platonists may think that there is a third term here -- an abstract object known as the Taller Than relation. But in the world of physical objects, Sam and Aaron are not joined by a third physical object that connects them in the Taller Than relation.

Hume's first idea set the stage for his second. If one event's causing another amounts to their standing in a certain abstract relationship, then it remains to say what this relationship consists in. Hume's regularity doctrine fills in the details here, claiming that there is nothing objective in the causal relation besides constant conjunction, contiguity, and temporal priority.

The regularity theory has been roundly criticized. Its successors typically introduce modal considerations as supplements or substitutes. Deterministic causality has been defined via different mixtures of the ideas of necessary and sufficient conditions. Probabilistic causality has been characterized in the same broad spirit. In the probabilistic theory, causes are not necessary for their effects; but they are necessary, in maximally specific background circumstances, to raise the probability of their effects.

Although these successors to Hume's theory part ways over what the causal relation amounts to, they usually

remain true to Hume's first ontological idea -- that causal connections are not third terms on an ontological par with causes and effects. In this paper, I'll begin with some intuitive reasons for doubting Hume on this first and fundamental point.

The proper reaction, however, will not be to jettison the search for modal concepts of causality but to supplement them with a conception in which causal relations are physical things. The context of argument, after a few everyday examples are explored, will be evolutionary theory. I will claim, not just that two concepts of cause are feasible in principle, but that each is useful in scientific practice. It may be of some interest that ordinary language permits us to talk in certain ways; but in matters of ontology, permission is of much less moment than necessity. Two concepts of cause, I want to argue, are needed in the framework of science.

2. The Kicked Golf Ball Meets the Sprayed Plant

The probabilistic theory of causality, which has been elaborated in slightly different ways by Good (1961-2), Suppes (1970), Cartwright (1979), Skyrms (1980), and Eells and Sober (1983) says, roughly, that positive causal factors raise the probability of their effects within maximally specific background contexts.² Critics of the theory, like Salmon (1980) and Otte (1981), have pinpointed two limitations in this proposal. First, one must assume that causal relations are indeterministic, if the relevant probabilities are to be well defined. Second, if one wishes to insure that an event's cause raises its probability, while other events with which it is correlated do not, then one must stipulate that causal chains form conjunctive, rather than interactive, forks, in the sense that Salmon (1978) has developed. Let us assume, therefore, that chance plays a role in the causal chains we are examining, where the forks in those chains are conjunctive, not interactive.

There is still at least one objection that needs to be addressed. Examples can be described in which it is plausible to claim that the cause lowers the probability of its effect. One such, briefly discussed in Eells and Sober (1983), is a modification of an example first suggested by Deborah Rosen. Imagine a golf ball rolling towards the cup, such that its probability of dropping in is quite high. Along comes a squirrel, who kicks the ball away from the cup, thus reducing its chance of going in. Then, through a series of improbable ricochets, the ball drops into the cup. How are the squirrel's kick and the ball's dropping in the cup related? I want to agree that the squirrel's kick reduced the probability of the ball's dropping in and that the squirrel's kick caused the ball to

drop in.

How, then, can I defend the probabilistic theory of causality? I do so by arguing that we must distinguish two concepts of cause. The first concerns the causal significance of properties, while the second addresses the causal significance of token events.³ We may wish to ask: in situations of the kind at hand, what is the significance of a squirrel's administering the kind of kick that the squirrel then produced? Or, one may ask, of that very event, what the squirrel's kick in fact produced.

At the level of properties, it is undeniable that the squirrel's kick was a negative causal factor for the ball's dropping in. The kick reduced the probability of the ball's going into the cup. It is possible, of course, that refined descriptions of the exact kind of kick the squirrel administered, of the position of the trees, etc., would reveal that this lowering of probabilities in fact does not occur. But, so as not to short circuit the example, let's assume that my claim that the squirrel's kick lowered the chance of the ball's going in the cup is ineliminable; it cannot be explained away. Evidence supporting this assessment could be assembled by replicating the set-up in a thousand control and a thousand experimental populations. At the level of kinds, it is right to deny that the squirrel's kick was a positive causal factor.

However, when we look at the token event, we recognize that the ball's dropping into the cup in some sense traces back to the squirrel's kick. The squirrel's kick produced the ball's trajectory. I assert that at the token level, causes don't have to raise the probability of their effects. What, then, is the nature of this causal relation between token events? To introduce this issue, I want to compare Rosen's example to one due to Cartwright (1979). What is curious about the next example is that it has the same probabilistic structure as the squirrel kick story. Yet, the stories are disanalogous when we consider token causal relations.

Cartwright (1979) considers an otherwise healthy plant that is sprayed with a defoliant. The defoliant lowers the plant's probability of surviving. Yet, improbably enough, the plant survives. The judgments made before about property causality remain in place. Before I said that the squirrel's kick was a negative causal factor with respect to the ball's dropping in the cup; now I say that spraying the plant was a negative causal factor with respect to the plant's surviving. Both are negative causal factors because both reduce the probability of the outcome.

However, shifting from property causality to token causality brings out a disanalogy. The squirrel's kick, I claim, caused the ball to drop in the cup. Yet, I think it is false that spraying the plant caused it to survive. The

examples are, or at least appear to be, identical with respect to probabilistic structure. This is why the examples are interchangeable when the question concerns property causality. Yet, they differ at the level of token causality. What, then, does this latter notion amount to?

Salmon (1978) has urged that the idea of a causal process is a useful one. I am not sure what sort of theoretical articulation it can be given. But if we allow ourselves to redescribe the two examples in terms of token processes producing outcomes, we can at least identify in a clearer way an asymmetry between the kicked ball and the sprayed plant.

The squirrel, when it kicks the golfball, initiates a new causal process, one that eventuates in the ball's dropping into the cup. The defoliant, on the other hand, does not produce a new process that causes the plant to remain alive. Rather, the plant remains alive because the processes initiated by the defoliant are arrested or prevented from getting through to the plant. The plant remains alive because the normal processes of physiological function, which were underway before the spraying, continued to operate.

Although I can offer no explication of this token causal relation, I will discuss one of its paradigm cases. This is the genealogical connection between parents and offspring. We use words like 'generates' and 'produces' to describe this relationship. Parental characteristics are often positive causal factors with respect to the characters of offspring. But this need not always be so. Imagine a simple haploid asexual population in which the probability of mutation is quite low. In such a population, like almost always produces like. We will consider a single locus at which there are two possible alleles, which I'll call '0' and '1'. Organisms may differ because they have genes at other loci that influence the probability of mutation at the locus of interest. And they may differ because they live in slightly different microenvironments that may influence the chance of mutation. But regardless of the genetic and environmental backgrounds, the probability that an offspring has the 1 allele is raised by its parent's having the 1 allele and lowered by its parent's having the 0 allele. The same holds for the 0 allele.

Let's consider a particular mother and daughter. The mother has character state 0 while the daughter, improbably enough, has 1. The parent's trait is a negative causal factor for the offspring's. Yet the mother generated or produced the daughter. The mother's character state was the token-cause of the daughter's. Or to put the idea a third way, the offspring's characteristic traces back to the parent's. Genealogy is a kind of causal connection distinct from property causation.

What is true of organisms is also true at higher levels of organization. We speak of species being related to each other as ancestor to descendent. Perhaps the characteristics of a parent species are usually positive causal factors for the characteristics of daughters. But this need not always be so. The ancestor's traits may have been negative causal factors for the descendent's. Nevertheless, the ancestor was the ancestor, and the token causal relation of "generation" connects them.

Another difference between property causality and token causality is worth noting. In saying that smoking is a positive causal factor for heart attacks in some population, we have in mind the way smoking affects each individual's chance of a coronary. But when an individual smoker has a heart attack, it is not inevitable that smoking did the causing in this token case. In a certain sense, the idea of property causality does not imply that any actual physical connection obtains between tokens of the two types. Smoking could be a positive causal factor for coronaries in a population even if no one smoked and even if smokers never had heart attacks. The claim of property causality could be true, I suggest, even if no smoker who gets a heart attack has his heart attack because he smoked.


When we think about token causal relations, we are considering an actual physical connection between cause and effect. It is with respect to this causal notion that Hume was wrong in his ontological claim. Besides the cause and the effect, there is the causal connection. For two events to be so linked is not merely for them to be relata of some abstract relationship. Rather, for one event to token-cause another is frequently for them to be connected by some third event or chain of such. The concept of a process may help in describing this actual physical connection. But I suspect that more is needed than Hume's ideas of contiguity, temporal priority, and constant conjunction, even when these are supplemented with modal considerations.

My reason for saying that causal connections are "frequently" third terms in the relation of cause to effect can be seen by considering an example discussed by Hume in the Treatise. A moving object (a billiard ball, say) collides with one that is at rest; the moving object stops and the stationary one goes into motion. Hume (1739, pp. 76-77) says that "when we consider these objects with the utmost attention, we find only that the one body approaches the other; and that the motion of it precedes that of the other, but without any sensible interval." I would suggest, however, that further attention reveals that the causal connection of these two events is mediated by a transfer of energy. This is the actual physical connection underwriting this case of token causality. The collision causes an energy transfer; that is how the first object's

motion caused the second to start to move. But suppose we now examine the connection of the collision and the energy transfer. Is this causal link underwritten by yet another physical process? To require that cause and effect are always mediated by physical connections is to stipulate in advance that there are no ultimate and unanalyzable token causal relations. But this we should not do: perhaps, at some level, cause produces effect, but not via the mediation of any third "physical process".

3. Probabilistic Causality and the Units of Selection

Before discussing the difference between selection processes that may occur at different levels, I'll say a little about how causality figures in the general idea of natural selection. It is standardly observed in evolutionary theory that a trait may become common because of a selection process without there being selection for that trait. This can happen if there are correlations of the right sort. The following figure shows two mechanisms that can make this happen. Arrows indicate causality and broken lines represent positive correlation.



Two phenotypic traits exhibit pleiotropy when they are joint effects of a single gene. One may be very advantageous while the other may be neutral or even slightly deleterious. Selection for the advantageous one can drive both to fixation (100% representation). The correlated trait is selected, but it is a "free rider"; there is no selection for having that neutral or deleterious character.

The mechanism of genetic linkage can achieve the same effect. Two genes may be close together on a chromosome. If one is very advantageous while the other is neutral or slightly deleterious, selection for the former may drive both to fixation. Again, the second character is selected (it free rides), but there is no selection for it.

When two characters are perfectly correlated, they are equal in fitness. This is because the fitness of a characteristic is simply the average fitness of the

organisms that have it. How, then, to represent the fact that two equally fit characters may yet differ in the advantages and disadvantages they confer?

The problem and a plausible answer to it are illustrated by a simple children's toy. The toy is a cylinder filled with balls of different sizes. Disks at different levels have holes of different sizes; a good shaking distributes the balls to their respective levels. The green balls get to the bottom because they are the smallest.

Let's think of this toy as a selection machine. The name of the game is getting to the bottom. A ball's chance of getting there is its fitness. Note that the green balls and the small balls have identical fitnesses, because the green balls are the small ones. When we shake the toy, we select the small balls and, therefore, the green ones. So if we talk about the fitness of characteristics and the selection of characteristics, there is no difference between being green and being small. A difference appears when we talk about selecting for this or that trait. There is selection for being small, but none for being green. The prepositions 'of' and 'for' mark an important distinction. 'Selection for' describes the causes of change; 'selection of' records the effects of certain processes.

When characters are correlated, either through pleiotropy, or gene linkage, or in the toy, how are we to disentangle their causal roles? One way is to consider how an organism's chance of reproductive success would be affected by each of the four possible combinations of the two characters considered. In the pleiotropy case, we want to consider how having both, one, or neither of E_1 and E_2 would matter. In the toy, we already see what happens to small green balls and to balls that are neither small nor green. But there are two, counterfactual, circumstances we also need to investigate. What would happen if a ball were green but not small? What would happen if it were small but not green?

The probabilistic theory of causation codifies this intuitive strategy. For being green to be a positive causal factor with respect to getting to the bottom, being green must raise the chance of getting to the bottom in every "background context". I put this last phrase in quotation marks, to indicate that it needs spelling out. I won't discuss this now, but will merely illustrate how it applies to the toy. Note that whatever color a ball may have, the ball will have a better chance of getting to the bottom if it is small than it will have if it is not small. So being small is a positive causal factor. But the same cannot be said about being green. One can't say that whatever size a ball may have, its chance of getting to the bottom will be enhanced by its being green. Given a ball's size, its chance of getting to the bottom is the

same, regardless of its color. So being green is neither a positive nor a negative causal factor.

The idea of selection for characteristics, I have suggested, needs to be understood causally. There is selection for a characteristic in a population precisely when having that characteristic is a positive causal factor for survival and reproduction. I also have suggested that the probabilistic theory of causality helps explain what 'positive causal factor' means in this context. Let me now indicate how these ideas can be applied to the units of selection controversy, particularly to the idea of group selection for altruism.

The standard, Darwinian view of natural selection pictures it as a process that occurs at the individual, organismic level. The rough idea is that individuals compete against each other in a single breeding population. Characteristics that promote the reproductive success of their possessors will increase in frequency; those that are less advantageous will tend to disappear. Let's consider two characteristics, which I'll call altruism and selfishness. Altruists donate reproductive advantages to others whereas selfish individuals may perhaps receive such benefits, but never give them away.

Continuing to speak somewhat roughly, we may say that the point of Darwinian, individual selection is that it is better to receive than to give. Under a process of individual selection, altruism will tend to disappear and selfishness will tend to become universal. And if, somehow, altruism were to become common in a population, it would be subject to "subversion from within". Sooner or later a mutant or migrant selfish individual would appear; that individual would be at a selective advantage and selfishness would go to fixation. According to this two-fold argument, altruism cannot evolve and cannot be maintained in a population. And this is so, despite the fact that a group of cooperating altruists would do better than a group of antagonistic selfish individuals. The good of the group or species is not what matters to Darwinian, individual, selection. It follows that if you observe "helping behavior" in nature, the Darwinian model of selection tells you that you should not try to explain it by positing a group selection process, one in which groups compete against other groups, where the trait that most benefits the group will ultimately prevail.

Some discussions of this problem, both in and out of biology, have tended to take the altruism out of altruism. This is the way that Trivers' (1971) has described his own idea of "reciprocal altruism". Reciprocal altruism is not altruism at all; rather, it makes explicit the intuitive idea that mutual cooperation can at times be more advantageous than antagonism. Suppose that the individuals in a population are somehow able to remember the behaviors

of others, and to reward and punish them accordingly. Although this kind of model may help explain some helping behaviors, we should recognize that it does not address altruism and selfishness in the sense just discussed. Consider three types of individuals C_1 , C_2 , and S . C_1 and C_2 cooperate with each other, but do not interact with S . C_1 and C_2 are each fitter than S as a result. But how are the fitnesses of C_1 and C_2 related? If they are equal, then C_1 and C_2 individuals do not, on average, sacrifice a reproductive advantage that benefits the individuals with which they interact. So neither type can be called "altruistic". If, on the other hand, C_1 and C_2 are unequal in fitness, then the fitter one will be driven to fixation by individual selection. Individuals with the fitter trait, on average, get more than they give in their interactions. So it is not altruistic. Altruism is not treated as a reality in Trivers' (1971) discussion, but is an appearance that must be explained away.


A second, individual-level, approach to altruism has used Hamilton's (1964) concept of inclusive fitness. Hamilton recognized that a gene may lever itself into the next generation in ways more devious than the obvious technique of making the organism in which it occurs have a lot of babies. If the gene can get the organism in which it is housed to help other organisms that also have copies of the gene, this may allow the gene to increase in frequency. But how can a gene in one organism make sure that benefits are donated to other bearers of the same gene? Hamilton's (1964) insight was that relatives have predictable probabilities of sharing genes. So, roughly, if a gene causes its bearer to help relatives in certain ways, this can allow the gene to increase in frequency.

This proposal is in the same spirit as Trivers' (1971); it may help explain helping behavior, but it does not explain altruism, properly understood. To see why, let us imagine a population containing two types. R individuals help relatives in ways that augment their inclusive fitness. S individuals do not. R will increase in frequency. Indeed, R will be a fitter trait than S . R is not altruistic. If individuals wish to insure that their traits are represented in the next generation, they would, quite selfishly, choose to be R .

We now can describe the concept of altruism more carefully. An altruistic characteristic must have two properties. First, regardless of the mix of altruists and selfish individuals found in a group, altruists, on average, do less well than selfish individuals. Second, a group of altruists does better than a group of selfish individuals. Although organisms in a group who behave the way Trivers (1971) or Hamilton (1964) describe may exhibit helping behavior, the helping behavior is not altruistic in this sense. The reason is that Trivers (1971) and Hamilton (1964) give models of individual selection that are

intended to explain the presence of helping behavior. However, altruism, in the sense I have described it, will be selected against at the individual level. Trivers (1971) was quite right to say that his idea takes the altruism out of altruism.

The fitness relationships required by the concept of altruism can be represented as follows:



\underline{S} and \underline{A} represent selfishness and altruism, respectively; \bar{W} represents the average fitness of organisms in the group. No matter what the mix is of \underline{S} and \underline{A} in a group, \underline{S} has a higher average *per capita* fitness. Yet, the average fitness (\bar{W}) in a group in which \underline{A} is common is higher than the average fitness in a group in which \underline{A} is rare.

This fitness relationship may be characterized in terms of the probabilistic theory of causality. There are two properties that influence the fitness of an individual. There is, first, the question of whether the individual is \underline{S} or \underline{A} ; second, there is the question of the frequency of \underline{S} and \underline{A} in the group that the individual inhabits. The fitness functions suggest¹⁰ the following two propositions: Regardless of what kind of group you are in, you'd be better off as a selfish individual than as an altruist. And regardless of what your individual phenotype is, you'd be better off in a group in which altruism is common than in a group in which it is rare. Selfishness is a positive causal factor, but so too is membership in an altruistic group. It is for this reason that individual, and group selection will oppose each other in this case.¹¹

Given that there is individual selection for selfishness and group selection for high concentrations of altruism, what will happen? This is an empirical matter, to be answered by consulting various contingent properties of population structure. Do individuals tend to live with each other irrespective of whether they are altruists or selfish, or is there a tendency for like to live with like? If kin live with kin, altruism may have a chance of evolving and being maintained in the ensemble of populations. Another relevant question is how severe the within group disadvantage of altruism is. A third concerns the rate at which new groups are founded and old ones go

extinct. The outcome of this process depends on these parameters; the simple observation that altruism is selected against within each group does not settle the question of whether altruism will be maintained or eliminated.¹²

The representation I have given of the controversy between group and individual selection concerns the relative importance of two sorts of causal factors that affect the reproductive success of organisms. There is an individual's phenotype on the one hand and its membership in a group of a certain phenotype on the other. Much of the dispute concerning "genetic selectionism" is in this mold (Sober and Lewontin 1982); it concerns the relative importance to an organism's reproductive success of its possessing a copy of this or that gene. That is, one dimension of the units of selection controversy can be viewed as focusing on organisms and asking what sorts of properties play a significant role in affecting their survival and reproductive success. I will say that this approach to the units of selection problem uses an organismic benchmark.

But other questions that have been prominent in the units of selection problem cannot be represented in this way. There are several possible selection processes that cannot be understood in terms of an organismic benchmark. I'll give two examples.

Meiotic drive is a process in which a gene in a diploid organism secures for itself greater than its "fair" 50% representation in the gamete pool. The normal process of gamete formation for a heterozygote with genotype Aa at a given locus will result in 50% A bearing and 50% a bearing gametes. But a driving gene will impair gamete formation in the gene it is next to. So in heterozygotes perhaps as many as 95% of the gametes will contain the driving gene. If this process were not counteracted by a force working in the opposite direction, such driving genes would go to fixation. In the house mouse and in fruitflies, there happens to be counteracting selection of this sort. But my interest here is in the process of meiotic drive, divorced from the other processes that may complicate the life of a driving gene.

There need be no differences in the fitnesses of organisms, if a driving gene, once it appears in a population, is to go to fixation. Organisms with copies of the gene may be exactly as viable and fertile as organisms without. Presence or absence of the driving gene makes a difference in gamete formation and in the genotypes of offspring. But having a copy of a driving gene need not be a positive or a negative causal factor for organisms. We are therefore compelled to interpret meiotic drive in a framework that differs somewhat from that suggested above. Many selection processes, including many that are discussed

under the rubric of genic selection, involve positive and negative causal factors pertaining to the reproduction of organisms. But meiotic drive requires that we abandon the organismic benchmark; here we must think of genes themselves as the objects possessing positive and negative causal factors for their own perpetuation.

At the other extreme in the hierarchy of levels of organization are species. Recently there has been interest in the evolutionary mechanism called "species selection".¹³ According to this idea, species in some lineages have characteristics that make them speciate more often. This process also requires that we depart from the initial framework in which it is organisms that possess the relevant causal factors. For two species may in principle differ in their propensities to speciate without their member organisms differing in their probabilities of survival and reproductive success. Imagine two populations. One grows larger while maintaining its internal integration; the other periodically fissions and forms isolates. The lineage produced by the latter may well end up containing more species, even though organisms in the two lineages do not differ in fitness. Here the positive and negative causal factors one wishes to examine attach to species. As in the case of meiotic drive, one must abandon the organismic benchmark.

Evolutionary theory now deploys a striking hierarchy of possible selection mechanisms. Indeed, the hierarchy is a double one. There is a set of questions concerning the various sorts of properties that may be significant causal factors in the life prospects of organisms. It is in this context that a number of debates concerning group selection, kin selection, and genic selection have been conducted. In addition, there are questions that force us to abandon the organismic benchmark -- questions in which the causal factors considered must be thought of as attaching to objects at other levels of organization. It is a matter of continuing empirical controversy which of these various processes have been important. In consequence, the questions surrounding the units of selection controversy provide a rich context in which philosophers can ask questions about the meaning of causal claims in science.


4. Token Causality and Phylogenetic Inference

In Section 2, I argued that the genealogical relationship of ancestor to descendent is a case of token-causality, to be distinguished from the relationship of property causality. I have not given an account of that token causal relation, but now I want to discuss an aspect of its epistemology. How may genealogies be reconstructed from data concerning the similarities and differences that

obtain among species? This problem, I think, offers some general insights into the question of when and why explanations that postulate a common cause of some observed events are preferable to explanations that postulate separate causes. Rather than describe the limitations of earlier approaches to this problem,¹⁴ I'll confine myself to elaborating an analogy between the following two questions: When is it reasonable to hypothesize that two species have a common ancestor? When is it reasonable to hypothesize that two events have a common cause?

First, I suggest that the token-cause relationship is transitive.¹⁵ In the genealogical case this is pretty clear. It is less transparent in general, but let that pass. My point in bringing it up is to urge the following consideration. If the ancestral relationship is transitive, then it is quite misguided to ask when two species have a common ancestor. Assuming that life on Earth originated just once, this question has the trivial answer: always. EVERY pair of species ultimately traces back to a common ancestor. The right problem to consider is this: how are we to tell when species A and B share a more recent common ancestor with each other than either does with some third species C?

Suppose we observe that species A and B have characteristic 1 while C has characteristic 0. We want to see whether this observation supports the phylogenetic grouping (AB)C better than it supports A(BC) or the "null" hypothesis that says that no pair of these species forms a group apart from the third. These three alternatives are shown in the following figure.



I assume that the common ancestor of all three species has the 0 state; that is, 0 is the ancestral and 1 the derived form of the character.

Each of the species considered has a probability of exhibiting the 0 or 1 state, if its ancestor has that or

the opposite state. With species A, we thereby associate two probabilities, one being its chance of having state 1 if its ancestor has that state, the other being its chance of having state 0 if its ancestor has that state. Besides species A, B, and C, and the root species that is ancestral to all of them, we also will want to consider a hypothetical ancestor D, which is invoked by two of the competing phylogenetic hypotheses to impose a grouping on the three taxa we've observed. Unlike A, B, C, and the root, the character state of D is unknown.

The three trees we are considering may be viewed as different assemblages of the four arrows shown in the lower right of the preceding figure. To evaluate the three phylogenetic hypotheses, we must decide how these branches are most plausibly assembled into a tree.

I propose that the likelihood of a tree be taken to measure the degree to which the observations support it. The likelihood of a tree is not its probability, but is the probability it confers on the observations. Given three assumptions, it can be shown that (AB)C is the best supported hypothesis of the three considered.¹⁶ These are:

- (i) intermediate probability values: all probabilities are between 0 and 1 noninclusive;
- (ii) probabilistic independence: the chance of an event on one branch of a tree is independent of what happens on others;
- (iii) Backwards Inequality: on each branch of the tree, $\text{Pr}(1 \rightarrow 1) > \text{Pr}(0 \rightarrow 1)$.


This last requirement says that the chance of a species having a given characteristic is greater if its ancestor already has that state than it would be if the ancestor had the opposite state. In the vocabulary of the probabilistic theory of causality, this means that 0 is a positive causal factor for 0 and 1 is a positive causal factor for 1. I will not describe why I think that evolution obeys the Backwards Inequality (on which see Sober 1983, 1984a), but will briefly consider what the meaning of this condition is for a general principle of the common cause.

Suppose you observe three events A, B, and C, with characteristics 1, 1, and 0, respectively. Suppose you know that all three trace back to a common cause, one whose character state is negative for A and B, but positive for C. Given this, it is better to interpolate a proximate common cause specific to A and B than it is to interpolate one joining B and C. Similarly, it is better to posit a common cause of A and B than to claim (as the "null" hypothesis does) that no pair has a common cause apart from the third.

It is worth noting that the common cause explanation here selected, which involves positing species D, does not say what character state D in fact has. The likelihood advantage of the common cause explanation does not depend on saying whether the character state of the common cause is positive or negative with respect to the character states of A and B. The token-causal relations are specified by the common cause explanation and these are enough to make it more likely than the alternatives.¹⁷


An example that doesn't concern phylogenetic inference may clarify the general features of this principle. Suppose we examine three rivers at the point that each flows into the ocean. Rivers A and B are polluted there, but C is clean. We assume that the three rivers ultimately trace back to a common source, at which point the water is clean. Can we infer anything about whether A and B have a common source that neither shares with C? The alternative hypotheses and the observations are shown below.

Some natural and fairly meagre assumptions suffice here. I assume that clean water upstream is a positive causal factor for clean water downstream, and that the same holds for polluted water.¹⁸ This means that if you observe a portion of a river in one state, the hypothesis that the river was in the same state upstream has a higher likelihood than the hypothesis that it was there in the opposite state.



These ideas allow us to conclude that (AB)C is better supported than A(BC) or the null grouping. Notice that (AB)C requires only a single change from clean to polluted water, whereas A(BC) requires at least two such changes. This is why the method of phylogenetic inference I've described is called 'parsimony'; it bids us minimize the number of parallel independent evolutionary events ("homoplasies").¹⁹

What if we had observed that A and B are clean whereas C is polluted? Assuming still that the source common to all three has clean water, the following figure represents our choices.



The assumptions made above no longer allow us to conclude that $(AB)C$ is the best supported hypothesis. Note that it is no longer true that $(AB)C$ requires fewer changes from clean to polluted. It demands that there have been at least one such, but so do the other alternatives.²⁰

The principle of the common cause does not group A and B together apart from C whenever events A and B share a property that C does not have. It is not simply that A and B match, but the relationship of their shared property to the character state of the cause common to all three, that is decisive. The probabilistic theory of causality allows us to express this requirement as follows: $(AB)C$ is better supported than the alternatives, if the property of the root event (i.e., the cause common to all three) is a negative causal factor with respect to the states of A and B , but positive with respect to the state of C .

5. Conclusion

I have described two concepts of cause that play a role in evolutionary theory. The first I have called 'property causality', otherwise known as 'the probabilistic theory of causality'. The idea is that a positive causal factor in a population raises the probability of its effect in every background context. This concept can be used to bring order to the multiplicity of models and arguments that constitute the units of selection controversy.

The second concept I have called 'token causality'. Its relata are not properties in a population, but token events. Although I tried to motivate it as a distinct concept, I provided no explication of it. Rather, I attempted to illustrate its use in the context of phylogenetic inference. Postulating common ancestors is a case of positing common (token) causes. Understood in this way, we can see how to formulate and justify a principle of the common cause. I say 'a principle' here because I grant that there may be others; the principle I have described may be less general than one would wish. It remains for future investigation to see whether this is so.

Notes

- ¹My thanks to Ellery Eells for useful comments.
- ²Skyrms (1980), Eells and Sober (1983), and Sober (1984c) discuss a weakened form of this criterion: a positive causal factor must raise the probability of its effect in at least one background context, and must not lower it in any other.
- ³In Eells and Sober (1983) and in Sober (1984c) these are called, respectively, population-level and individual-level causation. The distinction also is drawn by Good (1961-2), Skyrms (1980), and Rogers (1981).
- ⁴Although Anscombe's (1971) main goal is to show that causation does not require determinism, some of her remarks imply that there is a notion of causation distinct from what I have called property causation. Thus, she says (p. 67) that "causality consists in the derivativeness of an effect from its causes... . Effects derive from, arise out of, come of, their causes. For example, everyone will grant that physical parenthood is a causal relation. Here the derivation is material, by fission. Now analysis in terms of necessity or universality does not tell us of this derivedness of the effect; rather it forgets about that. For the necessity will be that of laws of nature; through it we shall be able to derive knowledge of the effect from knowledge of the cause, or vice versa, but that does not shew us the cause as source of the effect. Causation, then, is not to be identified with necessitation."
- ⁵A drawing of the toy is on p. 99 of Sober (1984c).
- ⁶The traits are to be understood in terms of their behavioral consequences, not in terms of their proximate etiology. This is why "altruistic" and "selfish" individuals do not have to have mental states. See Sober (1986) for further discussion.
- ⁷The expression is Dawkins' (1976).
- ⁸Here I ignore Hamilton's (1964) calculation of the trade-off between the amount of benefit the recipient receives, the amount of cost that donation imposes on the donor, and the coefficient of relationship of donor to recipient.
- ⁹The anthropomorphic language of the last few paragraphs is dispensable.
- ¹⁰A fine point forces me to say 'suggest' rather than 'imply' here. Strictly speaking, the fitness functions

describe the average fitnesses of the two types. This does not require that each individual in a group would have a higher fitness if he were selfish than he would if he were altruistic. See Sober (1984c) for further discussion.

¹¹The fitness function shown in the figure is a generalized form of the Prisoners' Dilemma. It also happens to represent the situation investigated by Axelrod (1984), in which cooperation may have a higher net payoff when one averages over an ensemble of interactions, even though it is disadvantageous as a strategy within each interaction.

¹²The evolution of altruism thus requires that the ensemble of populations exhibit a version of Simpson's paradox. See Sober (1984c) for discussion.

¹³The hypothesis has been defended by Gould and Eldredge (1977), Stanley (1979), and Vrba (1980). It is discussed in Sober (1984c).

¹⁴For criticisms of the approach to the Principle of the Common Cause taken by Reichenbach (1956), Salmon (1975, 1978), and van Fraassen (1980), see Sober (1984a).

¹⁵Although property causality is not in general transitive, Suppes (1970) and Eells and Sober (1983) describe sufficient conditions for its being transitive.

¹⁶The proof is given in Sober (1983) and (1984a).

¹⁷In addition, the meagre assumptions of the likelihood argument do not always permit one to assign character states to hypothetical ancestors.

¹⁸This does not mean, I should add, that water tends to retain its ancestral condition as it flows downstream. The probabilistic condition we are considering is retrospective, not prospective, as it were. That's why it's called the Backwards Inequality. See Sober (1983) and (1984a) for discussion.

¹⁹The division among biologists concerning the merits of parsimony as a method of phylogenetic inference is nicely represented by Felsenstein (1978) and Farris (1982).

²⁰This illustrates why symplesiomorphies -- matches with respect to ancestral character states -- are not interpreted by parsimony as evidence of phylogenetic relationship.

References

- Anscombe, E. (1971). "Causality and Determination." Delivered as Cambridge University Inaugural Lecture. (As reprinted in Causation and Conditionals. Edited by E. Sosa. Oxford: Oxford University Press, 1975. Pages 63-81.)
- Axelrod, R. (1984). The Evolution of Cooperation. New York: Basic Books.
- Cartwright, N. (1979). "Causal Laws and Effective Strategies." Nous 13: 419-437.
- Dawkins, R. (1976). The Selfish Gene. Oxford: Oxford University Press.
- Eells, E. and Sober, E. (1983). "Probabilistic Causality and the Question of Transitivity." Philosophy of Science 50: 35-57.
- Farris J. (1982). "The Logical Basis of Phylogenetic Analysis." In Advances in Cladistics. Edited by V. Funk and N. Platnick. New York: Columbia University Press. Pages 7-36. (As reprinted in Sober (1984b). Pages 675-702.)
- Felsenstein, J. (1978). "Cases in which Parsimony and Compatibility Methods will be Positively Misleading." Systematic Zoology 27: 401-410. (Reprinted in Sober (1984b). Pages 663-674.)
- Good, I.J. (1961-2). "A Causal Calculus I and II." The British Journal for the Philosophy of Science 11: 305-318; 12: 43-51; 13: 88.
- Gould, S. and Eldredge, N. (1977). "Punctuated Equilibria: The Tempo and Mode of Evolution Reconsidered." Paleobiology 3: 115-151.
- Hamilton, W.D. (1964). "The Genetical Theory of Social Behavior. I and II." Journal of Theoretical Biology 7: 1-52.
- Hume, D. (1739). A Treatise of Human Nature. London: John Noone. (As reprinted (ed.) L.A. Selby-Bigge. Oxford: Oxford University Press, 1968.)
- Otte R. (1981). "A Critique of Suppes' Theory of Probabilistic Causality." Synthese 48: 167-181.
- Reichenbach, H. (1956). The Direction of Time. Berkeley: University of California Press.
- Rogers, B. (1981). "Probabilistic Causality, Explanation, and Detection." Synthese 48: 201-223.
- Salmon, W. (1975). "Theoretical Explanation." In Explanation. Edited by S. Körner. Oxford: Blackwell. Pages 118-145.

- . (1978). "Why Ask 'Why?' - An Inquiry Concerning Scientific Explanation." Proceedings and Addresses of the American Philosophical Association 51: 683-705. (As reprinted in Hans Reichenbach: Logical Empiricist. Edited by W. Salmon. Dordrecht: Reidel. Pages 403-425.)
- . (1980). "Probabilistic Causality." Pacific Philosophical Quarterly 61: 59-74.
- Skyrms, B. (1980). Causal Necessity. New Haven: Yale University Press.
- Sober, E. and Lewontin, R. (1982). "Artifact, Cause, and Genic Selection." Philosophy of Science 49: 157-180. (Reprinted in Sober (1984b). Pages 210-231.)
- . (1983). "Parsimony in Systematics: Philosophical Issues." Annual Review of Ecology and Systematics 14: 335-358.
- . (1984a). "Common Cause Explanation." Philosophy of Science 51: 212-241.
- . (ed.). (1984b). Conceptual Issues in Evolutionary Biology: An Anthology. Cambridge: Bradford/MIT Press.
- . (1984c). The Nature of Selection: Evolutionary Theory in Philosophical Focus. Cambridge: Bradford/MIT Press.
- . (1986). "Methodological Behaviorism, Evolution, and Game Theory." Synthese. (forthcoming.)
- Stanley, S. (1979). Macroevolution: Pattern and Process. San Francisco: W.H. Freeman.
- Suppes, P. (1970). A Probabilistic Theory of Causality. Amsterdam: North Holland Publishing Co.
- Trivers, R. (1971). "The Evolution of Reciprocal Altruism." Quarterly Review of Biology 46: 35-57.
- van Fraassen, B.V. (1980). The Scientific Image. Oxford: Oxford University Press.
- Vrba, E. (1980). "Evolution, Species, and Fossils: How Does Life Evolve?" South African Journal of Science 76: 61-84.