Philosophy of Science Association

Probabilistic Causality: A Rejoinder to Ellery Eells

Author(s): John Dupre

Source: *Philosophy of Science*, Vol. 57, No. 4 (Dec., 1990), pp. 690-698

Published by: The University of Chicago Press on behalf of the Philosophy of Science Association

Stable URL: <http://www.jstor.org/stable/187770>

Accessed: 22/01/2009 07:09

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=ucpress>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.

Philosophy of Science Association and The University of Chicago Press are collaborating with JSTOR to digitize, preserve and extend access to *Philosophy of Science*.

DISCUSSION:
PROBABILISTIC CAUSALITY:
A REJOINDER TO ELLERY EELLS*

JOHN DUPRÉ

Department of Philosophy
Stanford University

In an earlier paper (Dupré 1984), I criticized a thesis sometimes defended by theorists of probabilistic causality, namely, that a probabilistic cause must raise the probability of its effect in every possible set of causally relevant background conditions (the “contextual unanimity thesis”). I also suggested that a more promising analysis of probabilistic causality might be sought in terms of statistical relevance in a fair sample. Ellery Eells (1987) has defended the contextual unanimity thesis against my objections, and also raised objections of his own to my positive claims. In this paper I defend and amplify both my objections to the contextual unanimity thesis and my constructive suggestion.

In a recent paper, Ellery Eells defends a thesis commonly held by theorists of probabilistic causality against some objections of mine.¹ The thesis in question is that to be a (probabilistic) cause a factor must increase—or at least not decrease—the probability of its effect in every possible set of background circumstances. Hereafter I shall refer to this as the contextual unanimity thesis (or more simply, “the unanimity thesis”). Eells also objects to some positive suggestions I made in my (1984). Here I would like to reply to some of these objections.

Eells directs at least three main criticisms at my views. First, he charges that my objections to the unanimity thesis depend on the failure to appreciate that probabilistic laws are a relation among three things: a causal factor, a probabilistic effect, and a particular population; not just between the first two of these. Second, I made the suggestion in my earlier paper that probabilistic causality would be better explicated in terms of a notion of statistical relevance in a fair sample. Eells objects to my appeal to the notion of a fair sample that it is insufficiently explicated (which I admit), and that in so far as it can be explicated, it leads to paradoxical and unacceptable results. And third, he suggests that my positive suggestions are relevant not to the explication of probabilistic causality at all, but

*Received May 1988; revised August 1988.

¹I refer to Eells (1987), commenting on Dupré (1984), which is in part a reply to Eells and Sober (1983). For further references see any of the above. All references to Eells in the text are to Eells (1987).

Philosophy of Science, 57 (1990) pp. 690–698.
Copyright © 1990 by the Philosophy of Science Association.

rather, if at all, to the problem of inference to causal hypotheses; or in other words, the concerns that I address are methodological rather than metaphysical.² Hence, presumably, my arguments are at cross purposes to the theorists I intend to address. In this reply I shall say something about each of these points.

It will help to begin by reviewing very briefly one of my previous objections to the unanimity thesis, the observation that it has some apparently very unintuitive consequences. Suppose that, on average, smoking raises your probability of developing lung cancer by a factor of 12. I say "on average" because exactly how smoking will affect an individual's chances of becoming ill will no doubt depend on many things: how much you smoke, how you smoke, and no doubt many peculiar features of your physiology and environment. If such were the case, it would seem correct to say that smoking caused lung cancer. Yet these circumstances would be quite consistent with the possibility that there should be some small minority of the population with a physical constitution such that smoking actually *reduced* their chance of developing lung cancer. In this situation, the unanimity theory implies that the statement "smoking causes lung cancer" is strictly speaking false, since the effect is non-unanimous. Moreover since we do not have, and from a practical point of view probably never could have, evidence to exclude the existence of such a fortunate minority, it appears that we really have no reason to believe that smoking causes lung cancer. The apologists for the tobacco industry have been vindicated by philosophy.

Eells replies, plausibly enough, that my example ignores the necessary relativization of probabilistic causal claims to particular populations. In the example above we need only say that in the population of individuals with the relevant physiological anomaly, smoking prevents cancer, for the rest, it causes cancer.³ For the total population, according to Eells, smoking is a "mixed" factor, which is to say that it is neither positive

²As Eells points out, my discussion of samples, rather than subpopulations, may have invited this misinterpretation. Nevertheless, samples are the appropriate entities for my purposes. The point of a *fair* sample is that it correctly (or closely) corresponds to the distribution of causal factors in the entire population. Since I conceive of causes as factors which, averaging over interactions with other factors, have a positive tendency to produce a particular effect, if a cause operates in the population, it will do so in a fair sample, even if, *pace* Eells, it does not so act in some subpopulations. That my appeal to samples suggests a convergence between methodology and metaphysics, I take to be a virtue of my approach.

³The situation described constitutes what Eells refers to as causal interaction, which he has discussed elsewhere in some detail (Eells 1986). The only modification that he proposes to the contextual unanimity theory to accommodate such cases is that interacting factors must be held fixed *even if* they are themselves neither causes nor preventatives of the effect under consideration. This assures the possibility of the response discussed and criticized in the present paper.

nor negative for cancer. It is, I suppose, something of an embarrassment to this reply that it entails that we never, or virtually never, know what the population is to which our causal claims apply; part of the motivation for my objection was the thought that it would be nice if we could make causal claims about antecedently identifiable populations. The unanimity thesis, on the contrary, entails that we cannot know the population to which a causal generalization applies until we have identified all the causal factors relevant to the production of the effect we are investigating. This seems to me an unduly cautious position.

However, rather than pursue these difficulties immediately, what I would like to emphasize at the outset is that my example was intended not solely to embarrass the unanimity theory, but rather to point to what strikes me as a major difficulty with the entire metaphysical picture that underlies that theory. In particular, what seems deeply questionable about the unanimity theory is the implicit assumption that there *is* some determinate set of factors, any maximal set of which will have a univocal tendency to produce or prevent an effect in which we are interested. This assumption is an essential part of what I characterized (in Dupré 1984) as covert determinism. This, at any rate, is my most fundamental disagreement with the theory. At the end of this paper I shall return briefly to this point.

I shall also have more to say below on the question of how probabilistic causal claims should be relativized to populations. (Eells's views on this topic are discussed in greater detail in Dupré and Cartwright (1988).) I might, however, emphasize here that my own position certainly is committed to such relativization though, as indicated above, relativization to rather different, and more readily identifiable, populations. But for now I shall turn to Eells's comments on the (admittedly sketchy) positive proposal I offered in my earlier paper, and Eells's objections to my appeal to "fair samples".

I certainly did not attempt anything like a complete explanation of the notion of a fair sample. In part this was because, in theory at least, I take the idea to be quite unproblematic: a fair sample is a sample that accurately represents the distribution of (relevant) properties in the population from which it is drawn; a biased sample fails in this regard. The simplest use of samples is just to measure such distributions, as of smokers in the U.S. or red balls in an urn. A more interesting case is their use in controlled experiments. In the first case the only problem in determining the fairness of the sample (apart from obvious questions about chance) is that of ensuring the causal independence of the method of selection and the property under investigation. In the second case there is also a problem about the causal independence of the effect being investigated and the method by which the cause is introduced. It does seem to me that while these problems are far from negligible, they are finite and generally tractable.

If one really wants to know whether smoking causes lung cancer, one should take large samples of infants, (randomly selected from the population, so there is good reason to believe that one has a *fair sample*) divide them into two equal groups, and force one group to smoke for the rest of their—no doubt abbreviated—lives. (See Giere 1984, p. 284). The cases that have been mainly debated by myself, Eells and others are precisely those in which such an approach is contingently unavailable (in the present case, for obvious ethical reasons). Part of my motive for thinking that even in these cases the notion of a fair sample would be analytically useful was the thought that a controlled experiment provides the paradigm case of having good evidence for the existence of a probabilistic causal relation. Of course, it is not trivial to apply this notion to cases where we are dealing with naturally occurring causes in uncontrolled populations, and I admit that I did not provide any sort of account in my earlier paper of how this should be done. However, Eells himself has carried out this task to my satisfaction if not to his, and I am indebted to him for doing so. His elaboration of my suggestion is as follows (p. 110). We select samples of smokers and nonsmokers from the general population in such a way that other known causal influences on the effect under investigation, in this case the occurrence of heart attacks, occur in these samples with the frequency with which they occur in the general population.⁴ We conclude that smoking (probabilistically) causes heart attacks just in case the frequency of heart attacks is higher in the sample of smokers than in the sample of nonsmokers. The underlying rationale, of course, is the same as that for a controlled experiment: we aim to compare samples in which other causes occur with equal frequency—equal, in fact, to the natural population frequency. In the latter case we have a method with a good chance of bringing this about; in the former we can do no better than to try to cook up samples with this feature.

Presumably the motivation for my proposal by analogy with controlled experiments will do little to impress Eells. Eells, in common with anyone else who thinks that the unanimity condition provides an explication of the concept of probabilistic causation (p. 111), is committed to the pessimistic position that a controlled experiment gives absolutely no infor-

⁴Suppose no smokers exercise. We must then compare samples both of which contain no exercisers. There is no immediate problem since smoking and exercising cannot interact, though we would clearly want to investigate possibilities such as that smoking caused heart attacks by preventing exercising.

As Eells obliquely suggests (p. 111, n. 3) my proposal has much in common with the theory of Giere (1984). One difference is that Giere, in discussing spurious correlations, suggests (pp. 293, 301) that we adjust the frequency of other factors in a control sample to match their frequencies in a test sample. My proposal is rather to match the frequency of other factors in the test sample to their frequency of occurrence in the general population. This seems preferable in principle because, at least in my view, we are ultimately interested in the impact of a factor on the population with its actual distribution of other factors.

mation about the causal propensities of a population.⁵ This is just because a controlled experiment tells us only about average effects. It tells us nothing about the various positive and perhaps negative effects that may contribute to that average. This strikes me as a disastrous divergence between metaphysics and methodology, a general issue to which I shall return below.

At any rate, Eells also offers a direct objection to my claim. He asks us to consider a situation in which the probability of suffering a heart attack depends on smoking and exercising in the following way: among exercisers, smokers are less prone to heart attacks than nonsmokers, whereas among nonexercisers the effect of smoking is reversed. Eells points out correctly (p. 110) that, according to my proposal, whether smoking causes or prevents heart attacks in such a situation will depend on the frequency of exercisers. I agree, but fail to see why anyone should be surprised by such a conclusion. Had we been blessed with a somewhat different biochemical constitution, smoking might have been a prophylactic against heart disease. I cannot see why our propensity to exercise should be less able than the details of our biochemistry to effect our causal susceptibilities. Thus I agree with Eells that we may be able to distinguish a subpopulation in which smoking prevents heart attacks. I just do not see why this should be taken to contradict the claim that in the population as a whole, smoking may cause heart attacks. For Eells a contradiction does arise, since whenever a factor is positive for an effect in one subpopulation, and negative in another, it is said to be mixed, and thereby neither positive nor negative in the whole population.⁶ This maneuver is essential for the defense of the unanimity theory. However, unless we take unanimity to be an intuitively necessary ingredient of probabilistic causality, a more pragmatically useful conception of relativization will allow that a cause be *both* positive for a population *and* negative for some subpopulation of it. (The example of the physiological condition under which smoking decreases the probability of lung cancer was intended to make just this possibility intuitively plausible.) Thus, to reiterate a point from my earlier paper, while it may often be very valuable to know that there is a subpopulation in which the causal influence of a factor is reversed,

⁵We do learn, I suppose, that it is not the case that *C* causes not-*E*, and that there is some population in which *C* causes *E*, though with no clue as to what this population might be. This does not strike me as exciting news.

⁶There is also something disturbingly arbitrary about the absolute distinction between a factor that reverses the direction of another causal factor, as in Eells's smoking and exercising case, and one which merely moderates or increases the strength of another. If one is really serious about unanimity, it seems to me, one should restrict one's causal claims to the (presumably very small) populations in which a factor has an influence not only in the same direction, but also of a uniform strength. This would, of course, greatly increase the counterintuitive quality of the proposal.

I wish to argue that this constitutes a supplementation, not a refutation, of our causal claim about the larger population.⁷

My answer to Eells's criticisms concerning the relativization of causal claims to particular populations should by now be fairly clear. As already admitted, my approach requires such relativization. A controlled experiment only gives reliable evidence about the causal properties of the population from which the samples studied are drawn. And more generally, any theory that averages across different auxiliary causal factors will be sensitive to the particular distribution of such factors in particular populations. Thus I am certainly in agreement with Eells's view that probabilistic laws should be defined only over specified populations. Where I disagree, as should by now be clear, is just on how narrowly such populations must be defined.

It is important to note that the extreme relativization required by the unanimity thesis presents overwhelming epistemological problems. If there are n factors causally relevant to the production of an effect, then unanimity requires a positive or neutral effect in each of the 2^n combinations of these factors. For many cases of interest it seems likely that 2^n will exceed, or be of comparable magnitude to, the population in question. The human population, for example, certainly one of the largest we are likely to deal with in biology (and in most cases much too broad a category for generalizations concerning humans) is of the order of 2^{31} . It seems highly likely that the probability of heart attacks might be affected by as many as 31 factors. If so, then the average size of the causally homogeneous classes would be about one; many would be empty; and it would be quite impossible to determine whether the unanimity condition obtained. It might well turn out, for a causal factor with complex and variable interactions with other factors, that the population in which that factor was operative could only be specified by complete enumeration.⁸

One's view of the preceding problem will turn, to some extent, on one's view of the nature of probability. It would be possible for someone who believed that probabilities were to be understood as limiting frequencies in theoretically infinite classes to maintain that the objections

⁷Another reason for insisting on the legitimacy of averaging is that there may be causes that can have opposite effects in entirely homogeneous populations, so that to eschew averaging will leave one with no way whatever of constructing appropriate causal statements. This possibility, which I take to be a significant one, is discussed in detail elsewhere (Dupré and Cartwright 1988), so I shall not pursue it here.

⁸Of course, as implied above, it cannot be assumed that all these sets of conditions will actually be realized in a given population. In cases where they are not, the defender of unanimity could either declare the causal truth about the population epistemically inaccessible, or decide that only exemplified sets of conditions mattered. The latter would have some paradoxical consequences. For example, it might turn out that the birth of a baby with a particular unique physical constitution would make it cease to be true that smoking caused lung cancer. Neither solution looks very encouraging for the unanimity theorist.

outlined above were “merely” practical. However, to defend the unanimity theory in this way would open up gaps between probability and probabilistic laws, and, again, between metaphysics and methodology, that seem to me intolerable.

This leads me to a disagreement Eells expresses with another claim made in my earlier paper (and also implied in my discussion, above, of the possible relativity of the causal upshot of smoking to the population frequency of exercising). There I suggested that we should recognize the possibility—indeed likelihood—that the laws of nature might turn out to depend on contingent facts about the populations to which they applied. Eells’s response to this suggestion seems very surprising. For he seems to concede that insofar as the laws in question apply only to specific populations my claim is true. But he then dismisses the suggestion as being true in only a “trivial sense” (p. 112), citing Hempel’s (1965) distinction between fundamental and derived laws, and asserting that my claim applies only to the latter. The implication is that laws referring to particular populations are necessarily merely derivative. But earlier in the paper, as I have mentioned, he states that “probabilistic causality is a relation between *three* things: a causal factor *C*, a probabilistic effect of it *E*, and a *population within which C* is a causal factor for *E*” (p. 107, original emphasis). And as he is at pains to emphasize, this is a claim about what probabilistic causality really *is*. Thus Eells’s theory leads inexorably to the conclusion that *all* probabilistic laws are merely derivative; or at least that there are no real probabilistic laws, as opposed to just probabilistic facts about particular populations. I could hardly have hoped for stronger support for my claim that the unanimity thesis is a kind of covert determinism.

I imagine that most theorists of probabilistic causality, and perhaps even Eells, would not wish to conceive of all probabilistic laws as *necessarily* relativized to particular (token) populations. I take it that the laws of quantum mechanics, assuming that these are irreducibly probabilistic, are not generally conceived in this way. But Eells and I perhaps agree that the laws of biology, anthropology, economics, sociology, etc. must be so relativized. It appears that Eells wants to claim that these laws are all derivative, and therefore scarcely worthy of being denominated “laws” (compare Smart’s (1963) claim about biological laws). But to claim that all the laws of these sciences are derivative is to embrace a very strong form of reductionism. That is an issue beyond the scope of this reply.⁹ However, it is an underlying commitment that needs to be made explicit in this debate.

This brings me to the last of the issues to which I wish to respond.

⁹The kind of reductionism in question is discussed and criticized in Dupré (1983, 1988).

Contrary to Eells's suggestion, I was thinking primarily about metaphysical issues rather than methodological ones (though see note 2, above, for a possible source of confusion). My metaphysical view is, on the other hand, partly motivated by the epistemological prejudice mentioned above: that the clearest idea we have of when we are confronting a probabilistic law (at least at the macroscopic level), is when we have a good controlled experiment. I am inclined towards a metaphysical view of *what* we are then confronting that is sufficient to explain why a controlled experiment of a suitable kind should reveal it, but does not assume any more than is necessary for such an explanation. I suggest this principle can be generalized in a traditional empiricist spirit: our metaphysical views should never be stronger than is necessary to make sense of what we take ourselves to know, and how we suppose we have come to know it.

This leads me to reiterate my most basic metaphysical disagreement with Eells and other proponents of unanimity theories: I see no reason to assume that there *is* any ultimate finest level of description for causal laws. The phenomena in which we are interested have many causes that we know of, and doubtless many that we don't. For all we know they may have infinitely many causes, though I do not propose to commit myself either way on that question.

It is pretty clear that the unanimity theory grossly violates such a principle. It is, in fact, largely irrelevant to understanding actual methods of causal inquiry. It fails to explain the utility of controlled experiments. And, recalling the discussion above, for a population size of N , the concept of a complete partition among n factors will have no application unless $N \gg 2^n$, a condition I see no reason to expect will generally be satisfied. As I mentioned, the pragmatic hopelessness of the unanimity theory does not preclude the possibility of defending it as a metaphysical thesis about causality; my point is just that perhaps it should.

In fact, I am inclined precisely to reverse Eells's view about which issues are methodological and which are metaphysical. It seems to me that partitioning should be seen as a pragmatic response to the unavailability of a metaphysically justifiable approach. The correct way of understanding partitioning, on this view, is not that one can thereby hope to achieve causally homogeneous samples—since we have no reason to suppose that there are such—but rather as the best methodology available for attempting to approximate the notion of a *fair* sample¹⁰ when one is not in a position to work with *random* samples.¹¹ One cannot use a ran-

¹⁰Strictly speaking, this requires additionally that the statistics from partitions be weighted according to their numerical size.

¹¹I do not, of course, contrary to the implication of Eells's discussion on pp. 112–113, deny that failure to consider an additional factor, whether in partitioning according to the unanimity theory or in constructing a fair sample, may lead to incorrect results. As Suppes

dom sample, because one is forced to resort to naturally occurring smokers rather than introducing smoking into a randomly selected sample of people. In this case the best we can do is to partition our data with respect to factors known or suspected to be causally relevant to the outcome in question, since we know that non-randomness in these respects will tend to distort our conclusions. How far it is reasonable to pursue such precautionary measures is a difficult and important methodological issue, but not, I am arguing, an issue of fundamental metaphysical consequence.

Picking up a familiar refrain from this debate, Eells concludes his paper by reiterating the thesis that "average effect is a sorry excuse for a causal concept". If so, perhaps it is God who owes us an apology.

REFERENCES

- Dupré, J. (1983), "The Disunity of Science", *Mind* 92: 321–346.
- . (1984), "Probabilistic Causality Emancipated", in P. A. French, T. E. Uehling Jr., and H. K. Wettstein (eds.), *Midwest Studies in Philosophy IX: Causation and Causal Theories*. Minneapolis: University of Minnesota Press, pp. 169–175.
- . (1988), "Materialism, Physicalism, and Scientism", *Philosophical Topics* 16: 31–56.
- , and Cartwright, N. (1988), "Probability and Causality: Why Hume and Indeterminism Don't Mix", *Noûs* 22: 521–536.
- Eells, E. (1986), "Probabilistic Causal Interaction", *Philosophy of Science* 53: 52–64.
- . (1987), "Probabilistic Causality: Reply to John Dupré", *Philosophy of Science* 54: 105–114.
- , and Sober, E. (1983), "Probabilistic Causality and the Question of Transitivity", *Philosophy of Science* 50: 35–57.
- Giere, R. N. (1984), *Understanding Scientific Reasoning*, 2nd ed. New York: Holt, Rinehart and Winston.
- Hempel, C. G. (1965), *Aspects of Scientific Explanation*. New York: The Free Press.
- Smart, J. J. C. (1963), *Philosophy and Scientific Realism*. London: Routledge and Kegan Paul.
- Suppes, P. (1984), *Probabilistic Metaphysics*. Oxford and New York: Blackwell.

(1984, p. 58) remarks: "it is a difficult and delicate matter to determine when no further causes can be identified." (Suppes 1984, pp. 55–59 offers a skeptical line of argument highly complementary to the present discussion.) When I advocated the rejection of a "completeness requirement" on the adequacy of probabilistic laws, I was making a metaphysical claim based on the views elaborated in both my earlier and the present text. The present discussion should at least make clearer why I disagree with Eells's remark (p. 112) that "even [my] proposal should require consideration of *all* causally relevant background factors."