

# Probabilistic Causality and Causal Generalizations\*

Daniel M. Hausman

Theorists of probabilistic causation have failed to distinguish between different tasks. One problem is to understand generalizations such as, “Smoking causes lung cancer,” “Seat belts save lives,” or “Just a spoon full of sugar helps the medicine go down.” Some causal generalizations, like the examples I have just given, are immediately practical. Other causal generalizations, such as those that are central in economics may be more theoretical. Whether immediately practical or not, causal generalizations are problematic, because the cause they purport to identify are not invariably accompanied by their effects. They are in this way irregular.

As philosophers such as John Stuart Mill (1843) and, more recently, John Mackie (1980) have shown, such irregularity does not rule out the possibility that the underlying causal relations are deterministic. If a cause is a conjunct in a minimal sufficient condition for its effect, then the effect may fail to accompany the cause whenever any of the other conjuncts are absent. But why believe that there are minimal sufficient conditions for lung cancer involving smoking or for demand increases involving price drops? Why not formulate a theory of the probabilistic causality that is expressed in causal generalizations?

The fundamentally indeterministic relations identified by contemporary physics also seem to call for a theory of probabilistic causality. For example, the collision of a neutron with a uranium 235 nucleus raises the probability that the nucleus will decay, but it does not raise the probability to one. Contemporary physics tells us that such decay probabilities cannot be explained by underlying deterministic relations. Though some philosophers would deny that these indeterministic relations

---

D.M. Hausman (✉)

Herbert A. Simon Professor, Department of Philosophy, University of Wisconsin-Madison,  
5197 Helen C. White Hall, 600 N. Park Street, Madison, WI 53706-1474  
e-mail: [dhausman@wisc.edu](mailto:dhausman@wisc.edu)

\*I owe a special debt to Ellery Eells and Elliott Sober, with whom I spent many hours discussing the issues raised in this essay, to Joonsung Kim, who devoted a chapter of his dissertation to criticism of an earlier version, and to Christopher Hitchcock for letting me appropriate so many of his ideas as well as for his criticisms of an earlier version of this essay. I am also grateful to Helen Beebee, Nancy Cartwright, Malcolm Forster, Huw Price, Peter Menzies, Charles Twardy, Jim Woodward, and John Worrall, who offered helpful comments on earlier drafts, to students in two seminars at the University of Wisconsin and to audiences at the University of Maryland, the London School of Economics, Duke University, Monash University, and Sydney University.

are causal (Papineau 1989; Woodward 1989; Hausman 1998, ch. 9), most who have addressed the question maintain that contemporary physics reveals that there are indeterministic causal relations.

Philosophers have hoped that a single theory of probabilistic causality would account for both causal generalizations and the indeterministic relations identified by contemporary physics, though they have rarely attempted to extend their account to the causal generalizations of the special sciences such as economics. I shall argue that the issues raised by causal generalizations are largely independent of metaphysical questions concerning probabilistic causality. This argument does *not* suppose that the causal relations that underlie a claim such as “Smoking causes lung cancer” are deterministic. Whether the underlying relations are deterministic or not does not bear on the question of whether smoking causes lung cancer. Metaphysical theories of probabilistic causation should not be expected to provide truth conditions for causal generalizations or to guide us concerning how to make use of them.

## Four Distinctions

Although this paper is mainly concerned to trace problems in theories of probabilistic causation to the mistaken assimilation of the issues raised by causal generalizations to those concerning indeterministic causal relations, there are other dimensions along which the objects of theories of probabilistic causality differ. Some of these are closely aligned with the difference between causal generalizations and claims about indeterministic causality, while others cut across this distinction. In particular, one should draw at least the following distinctions.

1. *Relevance versus role.* By the “role” or “bearing” of a causal factor, I mean whether a causal factor is positive or negative, or in some way “mixed” for some outcome. A poison and an antidote are both causally *relevant* to death, but their *roles* are opposite. A factor or variable  $X$  is causally *relevant* to another factor  $Y$  if it has *any* bearing on  $Y$ , positive, negative or mixed. Causal generalizations both in science and especially in daily life are typically concerned with role, rather than merely relevance. (They wouldn’t be practical otherwise!) One wants to know whether smoking increases the risk of lung cancer, not merely whether it is somehow relevant. One wants to know whether lowering the price of a commodity will increase the demand for it, not just whether price changes are somehow relevant.
2. *Variables versus values of variables.* Consider a continuous variable  $Y$  that measures an agent  $A$ ’s income and a second continuous variable  $Q$  that measures the quantity of chocolate  $A$  demands.  $Y$  is causally relevant to  $Q$ , and over some ranges of values, one can say that  $Y$  has a positive impact on  $Q$  in the sense that as  $Y$  increases, so does  $Q$ . So one can speak meaningfully of both the causal role and causal relevance of one variable for another (given values of the other variables and within some range of values). Causal generalizations within a science such as economics usually concern or derive from claims about the role and relevance of variables. A claim such as the law of demand, for ex-

- ample, follows from a more general functional relation postulated between the quantity demanded of a commodity, its price, income, and other variables. At the same time, one can also speak – as is more common in practical causal generalizations – of the causal role of the *value* of a variable *X* on some effect as a contrast between the effect of the given value and the effect of some other value of *X* that is pertinent in the context (see Hitchcock 1993, 1995, 1996). Smoking one pack of cigarettes a day increases the risk of lung cancer compared to not smoking any cigarettes, but it diminishes the risk compared to smoking two packs a day. The first contrast is of course usually the pertinent one. Practical causal generalizations typically concern relations among values of variables.
3. *Homogeneous versus heterogeneous circumstances.* Causal factors, as ordinarily conceived, have different consequences in different circumstances. Exposure to small pox will not cause the disease in those who have been inoculated. Suppose one had a list of all of the variables that are relevant to whether someone contracts lung cancer, other than *S* (the number of cigarettes smoked) and consequences of *S*. A generalization concerns causally homogeneous circumstances, if and only if these other variables have unchanging values. Practical causal generalizations (unlike claims concerning indeterministic causal relations within physics) typically concern causal relations when the values of other causally relevant variables are *not* unchanging. They concern causal relations in heterogeneous circumstances – that is, across some range of causally homogeneous circumstances.
4. *Types versus tokens.* Smoking can be a “type-level” cause of lung cancer – that is, it can tend to cause lung cancer – even when it does not actually do so. This difference between tendency and upshot cuts across the other distinctions.<sup>1</sup> In this paper I am concerned only with type relations – that is, with generalizations concerning causal tendencies. All of the other distinctions, including the distinction between causal relations in causally homogeneous circumstances and causal relations in heterogeneous circumstances are distinctions *among* causal generalizations, not distinctions between claims about actual causation and claims about causal tendencies.

If one simplifies and supposes that questions about causal relevance always arise with respect to variables rather than values of variables, then one can draw Table 1.

**Table 1** Distinctions among causal generalizations

	Causal relevance (of variables)	Causal role of variables	of values of variables
Heterogeneous contexts	Demand is a function of prices, incomes, and other things	The cancer risk increases with the amount one smokes	Seat belts save lives
Homogeneous contexts	In circumstances <i>K</i> , neutron bombardment is causally relevant to decay	<i>Ceteris paribus</i> the demand for <i>X</i> is a decreasing function of the price of <i>X</i>	In circumstances <i>K</i> , $\Pr(Y = y)/(X = x^*) > \Pr(Y = y)/X = x')$

Clarifying these distinctions cuts through problems that have plagued theories of probabilistic causality. Theories that address the metaphysical question of what it is for a variable to be probabilistically causally relevant to some outcome belong in the bottom left-hand cell, while a metaphysical theory of causal role (if there is such a thing) belongs in bottom middle or the bottom right-hand cell. Practical causal generalizations *presuppose* that there are causal relations of some sort – whether deterministic or indeterministic – when the circumstances are causally homogeneous. Their job is to provide guidance when one does not know what the other causally relevant variables are and what their values may be. They are generalizations across homogeneous contexts. The puzzle they present is to say what sort of generalizations they are, not what causation is.

Before addressing these puzzles, let us use these distinctions to clarify the form of the causal generalizations. For definiteness, let us focus exclusively on the top right-hand cell – that is, on causal generalizations about the bearing of values of variables. Since the causal role of variables depends on background circumstances, causal generalizations should be relativized to some population  $P$ . Furthermore, claims about causal role always contrast the effect of one value of the purported causal variable to the effect of another value. This is trivial in the case of dichotomous variables, but in the case of non-dichotomous variables, the contrast should be made explicit. So I shall take the canonical form of a causal generalization to be:

In population  $P$ ,  $X = x^*$  as compared to  $X = x'$  causes  $E$ .

In the case of dichotomous variables, this can be abbreviated as “In population  $P$ ,  $C$  causes  $E$ ,” with the contrast between the effect of  $C$  and  $\sim C$  understood. For practical purposes, one should require that the increase in the probability of  $E$  due to  $X = x^*$  (or  $C$ ) be substantial, but I shall leave this requirement implicit.

## The Irregularity of Causal Generalizations

According to warning labels on cigarettes, the Surgeon General has determined that smoking causes lung cancer. What does this mean? Since not everyone who smokes gets lung cancer, smoking by itself is not a deterministic cause of lung cancer. But one need not surrender the view that causation is a deterministic relation. Similarly the fact that causal relations among economic variables are not invariable has not led economists to abandon a deterministic view of causation. Following Mackie (1980, ch. 3) one might say that deterministic causes are INUS conditions for their effects. Suppose that smoking were an INUS condition for lung cancer, and that price changes were INUS conditions for changes in quantity demanded (Hoover 2001, ch. 2). Although not by itself sufficient, smoking would be a conjunct in one or more minimal sufficient conditions for lung cancer. Since there are presumably also minimal sufficient conditions for lung cancer that do not include smoking, smoking would not be necessary for lung cancer either. If some of the relevant causal relations

are not deterministic, then causes cannot be INUS conditions for their effects, but they can be conjuncts in minimal sufficient conditions that fix some objective chance of the effect occurring.

Let  $G$  be the other conjuncts in a minimal sufficient condition for lung cancer that includes smoking, and call the disjunction of the other minimal sufficient conditions for lung cancer " $H$ ." On a deterministic view of causation, smoking makes a difference to whether someone gets lung cancer only given the presence of  $G$  and the absence of  $H$ . Smoking is "necessary in the circumstances" – that is, necessary when none of the other minimal sufficient conditions for  $C$  are present, and sufficient when all of the other conjuncts in one or more of the minimal sufficient conditions including smoking are present. In the population as a whole, smoking has no single causal role. It causes lung cancer only in those individuals in whom just the right background conditions obtain. This account takes causation to be a three-place relation between a cause  $C$ , its effect  $E$ , and background conditions  $K$  in which  $C$  is necessary and sufficient for  $E$ . In a science such as economics, a good deal is known about the background conditions in which, for example, an increase in the money supply is necessary and sufficient for an increase in the rate of inflation, but the conditions cannot be completely specified.

Analyses such as these – whether deterministic or indeterministic – reveal a problem, which Wayne Davis calls "the background conditions problem" (1988, p. 133). The problem is that the third place in the causal relation means that causes are only causes when the conditions are "right." One can avoid introducing a third place in the causal relation by quantifying existentially:  $C$  causes  $E$  only if there exist background circumstances in which  $C$  is necessary and sufficient for  $E$  (and further conditions to insure the asymmetry of causation are met). But without saying more and without knowing whether in the actual circumstances  $C$  is necessary or sufficient for  $E$ , one is left with a very weak notion of causation. To say that smoking causes lung cancer is surely to say more than that there are some circumstances in which smoking is a conjunct in a minimal sufficient condition for lung cancer. Though people can rarely specify precisely what the other conjuncts in the minimal sufficient conditions including the particular cause are or what other minimal sufficient conditions there are for the given effect, they nevertheless usually know important facts about what background conditions must obtain for  $C$  to cause  $E$ . For example, without knowing exactly the conditions in which striking matches is necessary and sufficient for them to light, most people know that matches need to be dry.

Since people do not have detailed knowledge of the background conditions, why suppose that there *are* any minimal sufficient conditions for lung cancer or increases in the rate of inflation? Why should one believe that causation really is deterministic? Faith that causation is deterministic seems not only unjustified, but pointless as well, because it leaves one unable to say anything except that for some people smoking causes lung cancer, while for others it is irrelevant, and for still others it may prevent lung cancer. One might argue that only a dogmatic attachment to a deterministic theory of causation lends credibility to such a vague and unhelpful account. Why not focus directly on relations that people can know something about,

such as the non-deterministic manifest relation between smoking and lung cancer in the actual inhomogeneous circumstances in which people live, smoke, and die?

Patrick Suppes pursues very much this line of thought.

[A] mother says, "The child is frightened because of the thunder", or at another time, "The child is afraid of thunder". She does not mean that on each and every occasion that the child hears thunder, a state of fright ensues, but rather that there is a fairly high probability of its happening. . . .

It is easy to manufacture a large number of additional examples of ordinary causal language, which express causal relationships that are evidently probabilistic in character. One of the main reasons for this probabilistic character is . . . we do not explicitly state the boundary conditions or the limitations on the interaction between the events in question and other events that are not mentioned. . . . A complete causal analysis is far too complex and subtle, and not to the point for which ordinary talk is designed. (1970, pp. 7–8)

Although Suppes here emphasizes the supposed conformity of probabilistic causality to ordinary language,<sup>2</sup> rather than the avoidance of a metaphysical commitment to determinism, he is following the line of thought sketched in the previous paragraph. He suggests that one can avoid invoking unknown minimal sufficient conditions by developing a theory of probabilistic causality.

The Surgeon General obviously means to say more than that smoking has some probabilistic relevance to lung cancer, be it positive or negative, and monetarists are claiming more than that the money supply has some relevance to inflation. We are warned that smoking significantly *increases* the probability of lung cancer or that increasing the money supply will lead to a non-trivial increase in the rate of inflation. So Suppes attempts to formulate a theory of what I have called "causal role."<sup>3</sup> One can do so without referring to some set of unknown conjuncts in minimal sufficient conditions, by maintaining that smoking causes lung cancer only if the probability of lung cancer conditional on smoking is larger than the probability conditional on not smoking *regardless of the circumstances*. Instead of a three-place relation between *C*, *E* and background conditions *B*, perhaps one can get rid of the intangible third place in the relation and analyze causation in terms of a two-place relation of statistical relevance. With the additional stipulation that smoking precedes lung cancer, this is basically Suppes' definition of a *prima facie* cause (1970, p. 12). Theories of probabilistic causation attempt in this way to dodge the background conditions problem.

## Contextual Unanimity

But one cannot simply forget about the circumstances and other causal factors. It may be that smoking is not a positive cause of lung cancer even though  $\Pr(C/S) > \Pr(C/\sim S)$ . As R.A. Fisher postulated, some genetic common cause of smoking and lung cancer could explain the correlation. Neither is it necessary that  $\Pr(C/S) > \Pr(C/\sim S)$ . Some gene that makes people likely to smoke might impede lung cancer so that smoking and lung cancer are not positively correlated, even though smoking causes lung cancer. The probabilistic relations between smoking and lung cancer may be misleading.

The solution to these difficulties adopted by most theorists of probabilistic causation has been to require that smoking increase the probability of lung cancer within *all* cells of a partition formed by taking into account all the other causes of lung cancer that are not themselves effects of smoking.<sup>4</sup> The cells in such a partition are causally homogeneous background contexts, and the proposal can be restated as the requirement that causes increase the probability of their effects in all causally homogeneous background contexts. John Dupré has dubbed this the requirement of “contextual unanimity” (1984). Since theorists of probabilistic causality quantify over causally homogeneous background contexts, they avoid reintroducing a third place in the causal relation. Rather than relativizing causal claims to specific causally homogeneous contexts, theorists of probabilistic causation maintain that *C* is a positive cause of *E* (in some population *P*) if and only if *C* increases the probability of *E* in *every* causally homogeneous background circumstance in *P* (and some other condition is met that guarantees causal asymmetry, such as temporal priority of the cause).

Eells and Cartwright take analyses such as these to constitute a metaphysical theory of what causation is. In my view, in contrast, such theories are an amalgam of metaphysics and methodology. They combine a metaphysical theory of probabilistic causation within individual causally homogeneous background circumstances and a view of how causal generalizations generalize across such circumstances. Although the details vary, the implicit metaphysical view is that in some causally homogeneous circumstance, *C* is causally relevant to *E* if and only if it is probabilistically relevant to *E*, and *C* precedes *E*. The causal generalization, “*C* causes *E*” is then taken to maintain that *C* causes *E* within *every* causally homogeneous background circumstance.

I thus suggest that *contextual relativity has nothing to do with the metaphysics of causation*. The metaphysics in theories such as those of Suppes, Cartwright, Humphreys (1989), and Eells consists of the claim that causation consists of statistical relevance and temporal priority of the cause, given some particular value for each of the other causally relevant variables (that is, within individual causally homogeneous circumstances). Contextual unanimity figures instead in the attempt to explain how causal generalizations can be true and useful. In particular, contextual unanimity is the easiest way to avoid relativizing causation to particular contexts: if there are any circumstances (or “subpopulations” in Eells’ terminology – 1991, ch. 1), in which smoking does not increase the probability of lung cancer, then smoking is not a cause of lung cancer in the population as a whole. By insisting on contextual unanimity, one is thus able to say more than merely that smoking increases the probability of lung cancer in some circumstances, though it may lower it in others, and be irrelevant in still others.

Two additional considerations motivate the requirement of contextual unanimity and the unwillingness to relativize causal claims to particular background circumstances. First, contextual relativity makes it easier for theorists to convince themselves (erroneously) that their accounts of causal generalizations are part of a metaphysical theory of probabilistic causal relations. If instead one concluded smoking could be said to be a cause of lung cancer only with respect to some contexts and not with respect to others, then the truth of unrelativized causal



generalizations, such as the Surgeon General's, would depend on there being an implicit specification of background contexts. But such a specification of favored background contexts has no place in a theory in a metaphysical account of what causation is. By insisting on contextual unanimity, the difficulty vanishes: there is no need to justify zeroing in on some contexts and ignoring others.

Second, scientists do not know what the causally homogeneous background contexts are against which smoking may cause lung cancer, and even if these contexts were known, individuals do not know which cell of the partition they occupy. The Surgeon General needs to offer advice that is applicable to people who are not all in the same causal circumstances. If smoking had the *same* bearing on lung cancer in every cell of the relevant partition of other causal factors, then the Surgeon General could warn people about the risks of smoking without knowing anything about their particular circumstances.

The requirement of contextual unanimity is thus an attempt to *evade* the irregularity of causal generalizations. It only comes into play if one attempts to generalize across different causally homogeneous background circumstances. It has no relevance to claims concerning the causal relevance of  $X$  to  $Y$  or the causal bearing of a value of  $X$  on  $Y$  within a single causally homogeneous background circumstance, which is where all the metaphysical action, so to speak, lies.

But the irregularity of causal generalizations cannot be evaded. Since, as the INUS analysis reveals, there is no reason to expect that a deterministic cause  $C$  of  $E$  will have the same bearing on  $E$  in every causally homogeneous circumstance, why should one stipulate that probabilistic causes must satisfy contextual unanimity? If probabilistic causes are INUS conditions for some objective chance of their effects occurring, one would expect them to be as sensitive to the background circumstances as deterministic causes; and there seems to be no general reason to suppose that the objective chance of their effect occurring will be increased by the presence of the cause in every homogenous context.

There are more specific grounds to doubt contextual unanimity. Consider a question posed by John Dupré: Should one conclude that smoking does not cause lung cancer if it were discovered that some people have a rare physiological condition that causes them to contract lung cancer more often if they do *not* smoke (1984, p. 172)? Indeed, no hypothetical case is necessary: smoking does not *in fact* increase the probability of lung cancer in every causally homogeneous background situation. For example, in some people smoking causes fatal heart attacks rapidly enough that it tends to prevent lung cancer. Although smoking would be an undesirable way for these people to prevent lung cancer, it would do the job.

If one requires contextual unanimity, one thus has to deny that smoking is a positive cause of lung cancer, or one has to restrict the population to which the surgeon general's causal generalization is supposed to apply (Glennan 2002, p. 124). In just the same way, it turns out that seat belts don't save lives. Brushing one's teeth doesn't prevent tooth decay. Caffeine doesn't wake people up. Increases in the money supply don't spur inflation. Aspirins don't alleviate headaches. And, I suspect that a spoonful of sugar does not help this bitter medicine go down. In short just about every causal generalization turns out to be false, unless one radically restricts its scope.



Ellery Eells is willing to bite the bullet and to conclude that smoking does not cause lung cancer. He maintains that it is instead “causally mixed” for lung cancer (1991, p. 100). In some cells of the partition it increases the probability of lung cancer and in some cells it does not. Imposing a requirement of contextual unanimity implies that causes are in fact typically causally mixed for their purported effects. Given the contextual unanimity requirement, the only truthful causal generalization the Surgeon General can make about the consequences of smoking for Americans in general is that sometimes it causes smoking and sometimes it doesn’t. But the Surgeon General is neither uttering this useless truth nor is he falsely maintaining that smoking increases the probability of lung cancer in every causally homogeneous context. In failing to capture what claims such as the Surgeon General’s mean, theories such as Eells’ are unable to distinguish useful and apparently true generalizations such as “Smoking causes lung cancer” or “Seat belts save lives” from useless and apparently false generalizations such as “Vitamin C cures cancer.” Contextual unanimity is self-defeating in the analysis of causal generalizations and irrelevant to the metaphysics of indeterministic causation (see also Woodward 1989, p. 374).

One can try to save the contextual unanimity requirement by hedging causal generalizations or restricting their scope. Presumably there is some condition *H* in which it is true that seat belts invariably increase the probability of surviving crashes. But without knowing what *H* is or having any idea whether *H* will obtain in the event of an accident that might befall me, the true restricted generalization gives me no guidance concerning whether to wear my seat belt. To save the contextual unanimity analysis of probabilistic causal generalizations in this way is to make these generalizations useless.

So probabilistic theorists who insist on contextual unanimity are no more successful than the deterministic theorist in analyzing claims such as “smoking causes lung cancer.” *C* can be a cause of *E* even though its bearing on *E* differs in different causally homogeneous circumstances. To interpret the causal generalization “*C* causes *E* in population *P*” as maintaining that *C* increases the probability of *E* in every homogeneous circumstance in this population implies that causal generalizations are almost all false or else have such a narrow or unclear scope as to be useless. Some other way to generalize across contexts is needed.

The right way to interpret causal generalizations is, I think, basically John Dupré’s. Dupré’s idea (which is developed more precisely by Eells (1987, pp. 108–110) and especially by Hitchcock (1998, pp. 282–290) is that one should hold fixed the frequencies of all the other factors relevant to lung cancer (apart from smoking and its effects) at their frequency in the actual population and see whether, against this background, the conditional probability of lung cancer, given smoking, is larger than the conditional probability of lung cancer, given non-smoking. Of course, if one knew what the causally homogeneous circumstances were, the role of the causal factor in each of those circumstances, and which circumstances individuals were in, then there wouldn’t be any need to do any averaging. So one should resist interpreting Dupré as calling for people to *construct* these averages from more detailed knowledge. One can instead learn the average effect from comparing

outcomes in treatment and control groups in randomized experiments or by inferences from observed correlations. On Dupré's construal, the Surgeon General's claim aims to provide just the information needed to decide whether to smoke by people who do not know how their propensity to develop lung cancer differs from the population average.

Although the truth of causal generalizations may depend on the relative frequencies of different causally homogeneous contexts, Dupré is not reducing causation to mere correlation. On Dupré's and Hitchcock's view – at least as I understand it – the generalization, “In population  $P$ ,  $C$  is a significant cause of  $E$ ” is true if and only if *in an ideal randomized experiment* the frequency of  $E$  would be appreciably larger among subjects taken from  $P$  who are exposed to  $C$  than among subjects taken from  $P$  who are exposed to  $\sim C$ . Although in general one would expect a correlation between  $C$  and  $E$  in the population, a correlation is neither necessary nor sufficient for an “average effect.” A correlation is not sufficient, because it might reflect the fact that  $C$  and  $E$  are effects of a common cause. In such a case the existence of the correlation would not underwrite action to bring about  $C$ , and causal generalizations are, of course, supposed to guide action. One cannot prevent a storm by putting a barometer in a pressure chamber and thereby preventing its reading from falling. The existence of a correlation is not necessary, either. It could be that  $C$  causes an increase in the chance of  $E$  at the same time as some common cause counteracts this correlation. For example, if those who live in rural areas where other causes of lung cancer are absent are more likely to smoke, there might be no correlation between smoking and lung cancer, or even a negative correlation, even though both those living in rural areas and those living in urban areas are more likely to get lung cancer if they smoke.

Eells criticizes Dupré's proposal because it implies that whether smoking causes lung cancer depends on the actual frequencies of other factors. Change those frequencies, and smoking may cease to be a cause of lung cancer. In Eells' view, like Cartwright's, “smoking causes lung cancer” is supposed to be a causal law, which should not depend on the actual frequency of background conditions. So Dupré's account is “a sorry excuse for a causal concept” (Eells and Sober 1983, p. 54). Dupré agrees that his view makes laws depend on frequencies (1984, p. 173), but argues that this implication should be accepted. I disagree. Among other undesirable consequences, Dupré's view implies that people can change causal laws. Eells's critique of Dupré would be decisive, if the task were to formulate a metaphysical theory of indeterministic causation or to develop an account of probabilistic laws.

But both Eells and Dupré treat two tasks as one. Each sees that the other's theory is inadequate to the task with which each is mainly concerned. Dupré's theory is inadequate as a theory of indeterministic causation, which is what Eells is mainly concerned with. Eells' theory is inadequate as a theory of causal generalizations, which is what Dupré is mainly concerned with. The situation resembles that of two carpenters, one of whom mainly pounds nails, while the other more often screws things together. Both believe that a good carpenter needs only one tool. So the first uses only a modified hammer and the second only an odd screwdriver. The first points out how badly the screwdriver drives nails, while the second points out how badly the hammer turns screws.

What is at issue in theorizing about causal generalizations is causal irregularity. The operation of causal factors, whether deterministic or indeterministic, varies from context to context, and guidance is needed when the details concerning the contexts are not known. Theoretical work may focus on individual contexts or homogeneous contexts, because it need not necessarily provide such guidance. But if one hopes to offer advice to people who do not know which homogeneous context they are in, one has to generalize across contexts in which the effects of causal factors are not uniform. The point of the Surgeon General's claim is to provide information about the dangers of smoking to people who are in many different circumstances and who do not know which causally homogeneous context they are in. A generalization such as "smoking causes lung cancer" summarizes the qualitative "average effect," and it consequently depends not only on the cancer-causing propensities of smoking in causally homogeneous background contexts but also on the actual frequencies of the contexts.

Eells takes issue with this line of thought and argues that Dupré's account of irregular causation leads to mistaken advice. He writes,

[T]he question of whether smoking is a population-level cause of lung cancer will turn on the population frequency of that physiological condition, and in an unacceptable way .... For example, a person contemplating becoming a smoker, and trying to assess the health risks, should not be so concerned with the population frequency of that condition, but whether or not *he* has the condition. That is, the person should be concerned with which *subpopulation* he is a member of, the subpopulation of individuals with the condition (a population in which smoking is causally negative for lung cancer) or the subpopulation of individuals without the condition (a population in which smoking is causally positive for lung cancer. (1991, pp. 103–104)

If, as Eells imagines, one knows the causal bearing of smoking on lung cancer in subpopulations in which contextual unanimity holds *and* one can find out which subpopulation one is in, then one should make use of the more specific information. So, for example, if smoking raised the probability of lung cancer in men but lowered it in women, then the Surgeon General's claim about the effect in the population as a whole, even if true, would be misleading. Rather than averaging across contexts in which smoking is positive, negative, or neutral for lung cancer, we should focus on its causal bearing in the context or subpopulation in which we find ourselves.

Dupré's and Hitchcock's formulations may misleadingly suggest that one begins with knowledge of the relevant causal factors and thus with a complete partition into causally homogeneous background contexts. One then determines the quantitative bearing of smoking on lung cancer in each of these contexts and the frequency of each context and thereby calculates the average effect of smoking on lung cancer. If this were an accurate description of the problem, Eells would be right to maintain that we should focus on the bearing of smoking on lung cancer in the contexts in which we find ourselves rather than averaging (though if we knew the particular context, there would be no need for generalizations across contexts and hence no need to impose a contextual unanimity condition either).

But this is not an accurate description of the problem. Nobody knows all the causal factors that are relevant to whether somebody contracts lung cancer. There is

no way to calculate an average effect by summing over the effects in causally homogeneous circumstances weighted by their frequencies. Instead one can infer the average effect by means of experiment or critical examination of observed correlations. The point of the thesis that causal generalizations state average effects lies in justifying drawing causal conclusions from experiments and observations. These conclusions are important because agents who have no evidence about what subpopulations they belong to or concerning how the risks of lung cancer, for example, vary across different subpopulations can do no better than to rely on the average effect of smoking in the population as a whole.

Most theories of probabilistic causality fail to cope with the problems of causal generalizations, because these theories misconstrue the problems as calling for a metaphysical theory of probabilistic causality. They wind up either with metaphysical views that are hopeless as accounts of causal generalizations or with accounts of causal generalizations that are hopeless as metaphysical theories of causation. When considering claims such as “Seat belts save lives,” knowing that there are subpopulations in which  $C$  is a cause of  $E$  – whether deterministic or indeterministic – is not to the point. What one wants to know is the causal significance of  $C$  for  $E$  when it is already suspected that  $C$  is “causally mixed” for  $E$ . Causal generalizations are supposed to help out here. Some do and some do not.

## When Are Causal Generalizations True and Useful?

On the average-effect interpretation presented in the previous section, a causal generalization such as “In population  $P$ ,  $X = x^*$  as compared to  $X = x'$  causes  $E$ ” is true if and only if (a) in  $P$   $\Pr(E/X = x^*) > \Pr(E/X = x')$  and (b) the probability difference in (a) is due to the causal influence of  $X = x^*$  as compared to  $X = x'$  in some causally homogeneous circumstance occupied by members of population  $P$ . (a) and (b) give truth conditions for causal generalizations concerning populations occupying causally heterogeneous circumstances in terms of generalizations concerning causal relations obtaining within particular causally homogeneous background circumstances. Theorizing about these latter relations is a task for metaphysics. Since the task is to elucidate causal generalizations, rather than to clarify the nature of causal relations, one can help oneself to whatever theory of causation one prefers, provided that it preserves some link between causation and probabilities.

On this account (in contrast to [Hitchcock 2001](#), pp. 219–220), causal generalizations can be true, yet useless or even seriously misleading. Suppose, for example, that eating French fries causes heart attacks among men in some circumstances and prevents heart attacks in women in some circumstances. At the level of the whole population, it turns out that eating French fries consequently increases the risk of heart attacks among men and within the population as a whole, but eating French fries lowers the risk of heart attacks among women. If the Surgeon General knew these facts and then announced only that eating French fries makes

heart attacks more likely, he or she would be culpably misleading. But provided that the correlation really is a consequence of causal relations in causally homogeneous circumstances between eating French fries and heart attacks, this claim would be true.

This truth condition preserves and supports the intuition that part of the explanation for why some causal generalizations are useful is that they are true. But as just pointed out, a causal generalization can be true and misleading and even harmful to a great many people. If there were a dozen significant subpopulations in which the causal facts concerning the relationships between values of  $X$  and some effect  $E$  differed wildly, and everybody in the population  $P$  knew the facts about the subpopulations and knew which subpopulation he or she belonged to, then the average effect in  $P$  would be of little interest.<sup>5</sup>

There is, I believe, a great deal to be said about when practical causal generalizations are worth making, but little of philosophical interest. Clearly “ $C$  causes  $E$  in population  $P$ ” is more worthwhile when  $\Pr(E/C) - \Pr(E/\sim C)$  is large rather than small. It is more useful when  $E$  is more important. It is more useful when people are better able to bring about or to prevent  $C$ . True causal generalizations will in general be more useful than false generalizations, though falsehoods can, of course, sometimes have good consequences. For example, a mistaken causal claim that smoking causes acne could serve teenagers well by leading them to stop smoking and thereby to avoid heart attacks and lung cancer later in life. But the fact that falsehoods may do good is usually a bad reason for enunciating them. Finally, if the consequences of  $C$  for  $E$  differ appreciably over different subpopulations, then it can be harmful to generalize over the whole population. It is usually better to generalize concerning the narrowest populations for which the information is available.

It is also difficult to say much of philosophical interest about how individuals should change their behavior when they come to believe a causal generalization. One possibility, which has been developed carefully by Christopher Hitchcock (1998, 2001) is to idealize and suppose that the agent can estimate the subjective probability that he or she is located in each causally homogeneous background context and that causal generalizations provide the agent with knowledge of the difference that values of  $X$  make to the probability of  $E$  in each context. Knowing his or her own preferences, the agent can then choose the action that maximizes expected utility. More ordinary cases when the agent has little idea what the homogeneous contexts are or which he or she may be in can be modeled as cases in which the agent’s expectation of the effect of  $X = x^*$  on the chance of  $E$  will coincide with the average effect in Dupré’s and Hitchcock’s sense.<sup>6</sup>

I am skeptical about this approach because of the extreme idealizations it requires. What I prefer to say is simpler. Suppose that an individual agent  $A$  belongs to some population  $P$  for which it is true that  $C$  causes  $E$ , and that there is no narrower population to which  $A$  belongs for which there is any information concerning whether  $C$  causes  $E$ . Then  $A$  should regard actions that cause  $C$  as increasing the probability of  $E$  (in accordance with the generalization), unless  $A$  has some reason to believe that he or she belongs to some subpopulation of  $P$  in

which  $C$  does not cause  $E$ . (For example, even though smoking causes lung cancer, those on death row in Texas probably do not have to worry about contracting lung cancer if they smoke.) Since, by assumption, agents are seeking guidance concerning what to *do*, specifically causal information is crucial. What matters to agents are the consequences of acting and bringing  $C$  about or preventing  $C$ , not whether  $\Pr(E/C) > \Pr(E/\sim C)$ . So what is useful to agents are specifically causal generalizations, not claims about mere correlations.

Consider, for example, the generalization, "Seat belts save lives." In the population of drivers in the United States, the probability of surviving accidents if one wears seat belts is significantly larger than the probability of surviving if one does not wear them (though, of course, people can argue about how "significant" the difference is). No doubt those who wear seat belts are on average more conservative drivers, and so some of the correlation between seat belt use and survival could be due to this common cause. But this common cause does not explain the correlation between seat belt use and survival among those who are in particular classes of accidents, and our knowledge of the mechanics of accidents supports the claim that seat belts really do save lives. "Seat belts save lives" is a true causal generalization.

This generalization is, moreover, worth formulating and in general worth acting on. This is so, even though there are certain classes of unusual accidents in which one is *more* likely to die if one is wearing a seat belt. That means that in some subpopulations the correlation is reversed, and this reversed correlation is equally a causal matter. If agents knew in advance which class of accidents they would be in, then the facts about the average effect of wearing seat belts in the whole population of drivers would be irrelevant. But before accidents occur, when the decision about whether to buckle one's seat belt must be made, there is no basis to assign individuals to the subpopulation of drivers who will be in those rare accidents in which seat belts diminish the odds of survival. A great many people consequently wear their seat belts (as they rationally should), because they believe that the causal generalization about the whole population grounds an expectation that they will be less likely to be injured or killed if they wear their seat belts.

Consider one more example. What should one say about the generalization, "Entering a hospital for treatment makes people more likely to die." Here there is a very significant correlation. The probability of death in the near future is much higher among hospital patients than among those who are not in the hospital. The problem with this generalization is that the increased mortality is not an average *effect* of going into the hospital. Although hospitals do kill people, the main explanation for the correlation is, of course, a common cause that sends people to the hospital for treatment and then kills them. So though there is a correlation, this causal generalization is false and does not provide a reason not to enter a hospital. There may be subpopulations, however, in which the correlation between death and hospital treatment is due to the dangers hospitals pose. For people with minor ailments whose local hospitals are exceptionally poor, the causal generalization, "Among people in this area with minor ailments, entering the hospital for treatment makes one more likely to die", could be true and useful.

## Conclusions

To understand causal generalizations, one must understand how and why people generalize. Metaphysical theories of indeterministic causation need not trouble themselves with such questions. The metaphysical task is to clarify the causal relevance of variables within homogeneous contexts. Theorists of probabilistic causation tried to accomplish this task at the same time as they undertook to provide truth conditions for causal generalizations. They offered a probability increase and temporal priority view of causation within causally homogeneous background contexts, and they imposed a contextual unanimity condition to specify when causal generalizations are true. They then ran these two theories together into the view that causation is statistical relevance in all causally homogeneous circumstances (plus temporal priority of the cause). But the two theories should be pried apart. I have offered no assessment here of the view of indeterministic causation as statistical relevance plus temporal priority within a given causally homogeneous circumstance. Whatever one thinks of it, it is separate from the contextual unanimity account of causal generalizations, which I criticized.

The central point is that at least two theories are called for rather than one. In attempting to address at the same time all six of the cells in Table 1, near the beginning of this chapter, probabilistic theories of causation have wound up failing at their tasks. They offer no solution to the conundrums of practical causal generalizations, because they collapse in the typical case where the causal factors are mixed in the population as a whole, and one cannot specify in any non-trivial way subpopulations or circumstances in which contextual unanimity is satisfied. At the same time, they obfuscate and complicate indeterministic causation by focusing on problems that have little to do with the metaphysics of indeterministic causal relevance. There is no single relation of “probabilistic causality” manifested in quantum physics, everyday practical generalizations, and the causal claims of special sciences such as economics. The attempt to tackle both these problems in a single theory is a mistake and should be abandoned.

## Notes

<sup>1</sup> Christopher Hitchcock argues that the type-token distinction confounds two distinctions: the distinction between claims about causal tendencies versus causal accomplishments (“actual causation”) and the distinction between the scope of claims of both sort (2001, pp. 219–20).

<sup>2</sup> The claim to match ordinary usage is questionable. Notice that Suppes attributes to the mother the claim that causes make their effects highly probable, which conflicts with his own theory of causation as probability *increase*. In the same empiricist spirit, Wesley Salmon argues that one should focus on statistical rather than deterministic relations, because statistical relevance relations constitute the evidence for claims concerning irregular causal bearing (1984, pp. 184–185).

<sup>3</sup> One also needs at least a semi-quantitative theory to distinguish significant from unimportant causes, but I will ignore these problems in this essay. I believe contrasts similar to those I shall



discuss with respect to the theory of causal bearing play an essential role in providing a quantitative account of causal role.

<sup>4</sup> It is not easy to define the relevant partition precisely. For an early, influential, but flawed account, see Cartwright 1979, p. 26, and for criticisms of the details of her account, see Ray 1992, pp. 231–240 and Hausman 1998, p. 198. Such views abandon any attempt to offer a reductive analysis of causation in terms of probabilities. An alternative proposal defended by Brian Skyrms (1980) is to require only that causes increase the probability of their purported effects in some cells of the partition and that they never decrease the probability. This view is subject to the same criticisms as the requirement of contextual unanimity.

<sup>5</sup> An anti-drinking poster on a college campus proclaims, “Drinking causes AIDS”. On the average-effect view, this claim is probably true. Many people are inclined to judge it to be false on the grounds that drinking does not bear the right *kind* of causal connection to a disease such as AIDS. Unlike sharing needles, a shot of whiskey does not carry the virus. Although “Drinking causes AIDS” misleadingly suggests a certain kind of causal connection, there is no need to build these domain-specific details into the truth conditions.

<sup>6</sup> Given the idealizations in this approach, this requires that the subjective probability agents assign to their occupying any particular causally homogeneous background context match its actual frequency and that the subjective conditional probability of the effect given the value of the causal variable in each context match the objective probability.

## References

- Cartwright N (1979) Causal laws and effective strategies. *Noûs* 13. Reprinted and cited from Nancy Cartwright. *How the laws of physics lie*. Oxford University Press, Oxford, pp 21–43
- Davis W (1988) Probabilistic theories of causation. In: Fetzer J (ed) *Probability and causality: essays in honor of Wesley C. Salmon*. Reidel, Dordrecht, pp 133–160
- Dupré J (1984) Probabilistic causality emancipated. *Midwest Stud Philos* 9:169–175
- Eells E (1987) Probabilistic causality: reply to John Dupré. *Philos Sci* 54:105–114
- Eells E (1991) *Probabilistic causality*. Cambridge University Press, Cambridge
- Eells E, Elliott S (1983) Probabilistic causality and the question of transitivity. *Philos Sci* 50:35–57
- Glennan S (2002) Contextual unanimity and the units of selection problem. *Philos Sci* 69:118–137
- Hausman D (1998) *Causal asymmetries*. Cambridge University Press, Cambridge
- Hitchcock C (1993) A generalized probabilistic theory of causal relevance. *Synthese* 97:335–364
- Hitchcock C (1995) The mishap at Reichenbach falls: singular vs general causation. *Philos Stud* 73:257–291
- Hitchcock C (1996) Farewell to binary causation. *Can J Philos* 26:267–282
- Hitchcock C (1998) Causal knowledge: that great guide of human life. *Commun Cogn* 31:271–296
- Hitchcock C (2001) Causal generalizations and good advice. *Monist* 84:218–241
- Hoover K (2001) *Causality in macroeconomics*. Cambridge University Press, Cambridge
- Humphreys P (1989) *The chances of explanation*. Princeton University Press, Princeton, NJ
- Mackie J (1980) *The cement of the universe*. Oxford University Press, Oxford
- Mill JS (1843) *A system of logic*. Longmans, Rpt. London 1949
- Papineau D (1989) Pure, mixed and spurious probabilities and the significance for a reductionist theory of causation. In: Kitcher P, Salmon W (eds) *Scientific explanation*. University of Minnesota Press, Minneapolis. *Minnesota Studies in the Philosophy of Science*, 13, pp 307–348
- Ray G (1992) Probabilistic causality reexamined. *Erkenntnis* 36:219–244
- Salmon W (1984) *Scientific explanation and the causal structure of the world*. Princeton University Press, Princeton, NJ
- Skyrms B (1980) *Causal necessity: a pragmatic investigation of the necessity of laws*. Yale University Press, New Haven, CT

Suppes P (1970) A probabilistic theory of causality. North Holland, Amsterdam

Woodward J (1989) The causal mechanical model of explanation. In: Kitcher P, Salmon WC (eds) Minnesota studies in the philosophy of science: vol. 8. Scientific Explanation, University of Minnesota Press, Minneapolis, pp 357–383

The Place of Probability in Science  
In Honor of Ellery Eells (1953-2006)

Eells, E.; Fetzer, J.H. (Eds.)

2010, XLIV, 368 p., Hardcover

ISBN: 978-90-481-3614-8