5

Probabilities for Explanation

D. H. MELLOR

I agree with Professor Salmon that the explanations he cites are causal and that they can dispense with determinism. Causal explanations of phenomena need only make them more or less probable, even if at some other (e.g., microscopic) level the phenomena are also in fact determined by "hidden variables." Causal explanation is thus a special kind of statistical explanation, deductive–nomological explanation being simply an extreme case that gives the explained phenomenon a probability of 1. The dearth of deterministic laws in the human sciences generally, and in archaeology in particular, makes this extreme case of little practical interest. And excessive emphasis on it by the "new archaeologists" has only served to discredit the idea of archaeological explanation as being either causal or lawlike. In fact it is both, although the laws may only be statistical, and it is not necessary to specify them in giving causal explanations, so long as the explanation clearly gives the phenomenon to be explained whatever probability is required.

There is of course more to causal explanation than probability, but I shall not discuss the other requirements, e.g., that causes precede their effects, nor why they matter. I wish only to make two points about the probabilities involved: one about the *kind* of probability, the other about the *level* of probability, that causal explanation demands.

Explanatory probabilities must be objective. There is a subjective theory of probability, which treats it merely as a measure of how confident people are of

THEORY AND EXPLANATION
IN ARCHAEOLOGY

ISBN 0-12-586960-6

whatever they think probable, but that theory is clearly inadequate to this application. It is not the business of explanation to make us more, or less, confident that the phenomenon to be explained has occurred. Usually we were already quite confident about that, and in looking for an explanation we are not looking to have our degree of confidence in it increased, or reduced, or confirmed. In explanation we are not concerned, as subjective probability is, with establishing to our satisfaction that the phenomenon to be explained has actually occurred. The idea that we are is the widespread misconception of explanation as a kind of argument or inference, a misconception that Salmon has notably and rightly combated.

By the same token, however, not all objective probabilities are explanatory. In particular, merely inductive probabilities are not. However objective they may be, merely inductive probabilities are still only measures of confidence: not of how confident we would be, but how confident we *ought* to be, of the occurrence of the phenomenon to be explained if all we knew was the proposed explanation of it. That is not at all what explanation is about.

The quickest way to see how merely inductive probabilities differ from the probabilities causation needs, without going too far into current philosophical controversy about probability, is by means of a simple instance of probabilistic decision making. Suppose the statistics suggested probabilities of .4 and .1, respectively, for smokers and nonsmokers getting cancer, and that (in some suitable units) I attach relative utilities of 10, 9, −10, and −11, respectively, to not getting cancer with and without smoking, and to getting it with and without smoking. This means in particular that I would prefer to smoke whether I had cancer or not. So the so-called "dominance principle" of decision theory tells me to carry on smoking. But the principle of maximizing my expected utility tells me to quit, since my expected utility if I smoke is 2 [$10 \times .6 - 10 \times .4$] and if I don't it is 7 [$9 \times .9 - 11 \times .1$]. Which principle should I follow—should I quit, or not?

It all depends on whether the probabilities involved are merely inductive or are what I shall call "causal probabilities." If whether I smoke or not actually *affects my prospects* of getting cancer, the probabilities are causal, and I should quit. But smoking might, for example, be just a statistical symptom of a genetic predisposition to get cancer whether I smoke or not, just as a falling barometer is only a symptom of ensuing rain. All the higher probability of cancer among smokers would show then is that smoking is some evidence for the presence of this predisposition. That is, the probability would be a perfectly good inductive probability, but my giving up smoking would not affect my prospects of getting cancer in the slightest—any more than putting a falling barometer under pressure will stave off rain. So, since I prefer to smoke in any case, I should follow the dominance principle and carry on.

I will *define* causal, as opposed to merely inductive, probabilities as *the objective probabilities that make the principle of maximizing expected utility rational to follow even when it conflicts with the dominance principle.* (I say "as opposed to *merely* inductive" because causal probabilities are inductive all right, i.e., they are a good measure of the confidence one should have in the occurrence of whatever has them. But that is not all they are.) Now it is obvious in the smoking example that probabilities must have this property if they are to support causal explanations. If smoking were only inductive evidence for the presence of incipient cancer, a man's smoking would no more cause or causally explain his getting the disease than a falling barometer causes or causally explains the subsequent rain for which it likewise is inductive evidence. Only if the probabilities of cancer among smokers and nonsmokers are causal in the sense just defined can a smoker's habit properly be said to cause or causally explain his getting cancer.

Now when searching for probabilistic causal explanations, one cannot expect to read the relevant probabilities straight off the statistics, in this case the proportions of actual smokers and nonsmokers getting this or that disease in some limited population. A correlation might be a freak statistical accident in a small population. Even in a large one, the statistics might just reflect inductive probabilities. Or they might be a mixture of inductive and causal probabilities, with different sorts of smokers having different causal probabilities of getting cancer. Here, as in the interpretation of statistics in archaeological examples, it is by no means straightforward to extract whatever causal probabilities there may be. Basically, it demands a theory to suggest possible causal mechanisms and means whereby their presence or absence can be tested for. That is a complex subject that I cannot embark on here: I only want to emphasize that it takes more than a statistical correlation to show the presence of causal probability. What more it takes is well illustrated in the successive reports on smoking and health put out by the Royal College of Physicians. The tobacco companies' objections to a causal interpretation of the statistics are certainly untenable now; but while there was only a correlation to go on, they had a case.

As in this simple and familiar case, so mutatis mutandis in the archaeological and other cases cited by Salmon. The explanatory force of the statistics appealed to rests in every case on the probabilities involved being causal rather than merely inductive, as a little reflection will reveal. Take the climatic explanation of Pleistocene extinction. The relative probabilities of survival of species in different climates must be taken to be causal for this to be an explanation. That is, by definition, it must be rational for an agent with a sufficient interest in wiping the species out and able to control the climate to alter it to do so, even though he would himself prefer a milder climate, whether the species survived or not. It must not for example be that extinction and climate are alike merely statistical effects of some independent cause, as rain and a falling barometer are of a lowering of atmospheric pressure. For then there would be no more point in altering the climate to wipe out the species than there is in trying to prevent rain by pressurizing a barometer; so the altered climate would no more explain species extinction than a rising barometer explains sunshine.

Which brings me to the second point I want to make—about the *level* of causal probability required for causal explanation. The phenomenon to be explained must be made more probable by the factors cited to explain it than in the circumstances it would have been had they not obtained. I need to make this point and to defend it because in the course of developing his statistical-relevance model of explanation Salmon was led to deny it, and the increasing influence of his model among archaeologists may mislead them on the matter. In particular, since the point is both true and important, I should not like Salmon's denial of it to discredit probabilistic theories of causal explanation and inhibit their acceptance by theoretical archaeologists.

Now I do not dispute the considerations that led Salmon (e.g., Salmon *et al.* 1971) to deny that causes must raise the probability of their effects. I merely think he was misled by them. First, he rightly observed that Hempel's reason for demanding high probability in statistical explanation—namely that high probability of the conclusion is a self-evident virtue in statistical inference— is no reason at all, because explanation is not a kind of inference. Just because the phenomenon to be explained would be more safely predictable if it were more probable doesn't mean it would therefore also be better explained.

But disposing of a bad reason for requiring causes to raise their effects' probabilities does not dispose of the requirement itself. Salmon, however, also has a positive reason for denying it. This is his observation that probabilities used in causal explanation must be relative to all statistically relevant causal factors (at the level of explanation involved—excluding for instance possible microscopic causes). A statistically relevant factor is one that affects the causal probability either way, whether raising or lowering it, and of course Salmon is right to insist that causal probabilities must be relative to all such factors. That is a precondition of the probabilities concerned being causal in the first place, rather than merely inductive. It is no use trying to raise the probability, and so beef up the explanation, of someone's getting cancer by suppressing the fact that he never smoked.

But this precondition, that causal probabilities must be relative to all relevant causal factors, does not mean that all causal probabilities provide equally good causal explanations. If it did, it would mean that a probabilistic causal explanation of a phenomenon occurring would serve just as well to explain its nonoccurrence, which is absurd. Someone's smoking would have to be as good an explanation of his not getting cancer as it is of his getting it, which it is not. Climatic changes that make the survival of a species less likely would have to be as good explanations of its survival as of its extinction, which they aren't. Take any of the examples Salmon cites, or any others seriously considered in archaeology or anywhere else. To practitioners not already brainwashed by the statistical–relevance model it will be quite plain that their acceptability as explanations depends on the causal factors they cite being not only statistically relevant but (in Salmon's terminology) positively relevant.

I suppose Salmon might protest here that our intuitions about this have been distorted by years of exposure to the Hempelian misconception of explanation as inference and need to be retrained. I doubt it: I reckon the intuition that causes need to raise their effects' probabilities predates Hempel and will be found flourishing among scientists who have never exposed themselves to contamination by the philosophy of science. But however that may be, Salmon would also remark at this stage that since improbable events are bound to occur occasionally, and their low causal probability nonetheless provides as good an explanation of them as there is, my requirement condemns some events to being absolutely inexplicable. So it does, and so they are . Of course, the fact that very improbable events occur from time to time is entirely explicable, since it is very probable that they will. But to the extent that each individual event is improbable, its occurrence remains unexplained. The fact that the proposed explanation is as good as there is does not mean that it is good enough, i.e., good enough to satisfy the original desire to know why the event occurred rather than not. Making hell inescapable does not make it heaven, and I see no a priori reason to expect the universe to scratch our every itch to understand it.

It may be tempting to conclude at this point that Salmon and I are quibbling as to which of two marginally different concepts has the better right to the title *explanation*—a terminological squabble of no serious interest to archaeologists or anyone else. The temptation should, however, be resisted for three reasons, of which the second especially has immediate application to a proper understanding of archaeological method. And although the first may seem rather abstract, it is worth mentioning because it helps to explain what may have appeared a bizarre invocation of decision theory to characterize causal probabilities.

What makes decision theory pertinent is that human action is the very paradigm of causation: Whatever else causation is, it is the mechanism that links the bodily movements we directly perform to other events and states of affairs we act, for whatever reasons, to influence or bring about. I do not mean that causation depends on human action—there could be plenty of causation in a world wholly devoid of agency—merely that our concept of causation has inter alia to make sense of action, including human action, as a phenomenon of the natural world. And that means in particular that it must always make sense to act in order to bring about an effect of that action. But only if causes make their effects more probable than in the circumstances they would otherwise have been will that be so, since otherwise the action would not make the effect desired more likely to occur.

This principle may be put in decision theory terms as follows. The effect of an action is the outcome of it that, if most desired, makes that action maximize expected utility when the action itself has no intrinsic utility, i.e., when the utilities of the outcomes are the same whether the action is performed or not, so that it must be done entirely for the sake of its effects. Suppose for example that smoking as an activity has no intrinsic utility for me in this sense, so that I must base my decision whether to go on smoking entirely on the utilities I attach to outcomes such as getting or not getting cancer. Then because as things are the probability of cancer is greater if I smoke than if I don't, maximizing my expected utility in this case means smoking if I would prefer to get cancer, and not smoking if I prefer not to. That makes cancer an effect of smoking and noncancer an effect of not smoking, and it does so because my smoking raises the causal probability of my getting cancer above the level it would have if I didn't smoke.

That is the fundamental reason causes have to raise their effects' probabilities, to make sense of causation's being the mechanism of action. But there is another reason as well, of more immediate interest to archaeologists, namely to make sense of their applications of the principle of inference to the best explanation. This is the principle that data that could be explained by several incompatible hypotheses give reason (*ceteris paribus*) to accept the hypothesis that would if true explain the data best. I cannot go here into the rationale, force, and limitations of this principle, but appeals to it occur all the time in everyday life, in law, in history, and in sciences and technologies, both natural and human. Archaeologists especially should be aware of it, since it provides their main reason for being interested in explanation in the first place. Very little of the archaeological record, after all, is interesting enough to be worth explaining for its own sake. Most archaeologists are prehistorians, interested in the record chiefly as evidence for the prehistoric happenings that caused it. They look for good prehistoric explanations of it because the better the explanation would be if it were true, the more likely it is to be true, other things being equal.

But how, in order to apply this principle, does one assess the comparative quality of putative causal explanations? Clearly there are many factors—simplicity, for one—that I cannot go into here. But there is also the causal probability that a proposed prehistoric explanation would, if true, bestow on the archaeological data it explains. The higher that is, other things being more or less equal, the more likely the explanation is to be the true one, i.e., the higher its *inductive* probability relative to the data it purports to explain. That is the so-called "maximum likelihood" principle of statistical inference. Now, is this principle really entirely independent of the principle of inference to the best explanation, as it would have to be if the level of causal probability supplied by causal explanations made no difference to their quality as explanations? If so, assessments of prehistoric hypotheses as prospective explanations, and by how probable they would if true make the data, must be entirely separate processes leading, if they conflict, to a further process of adjudication between them. And obviously they are nothing of the sort. Comparing the causal probabilities rival prehistoric hypotheses invest the archaeological record with is all part of finding out how good they would be as explanations if they were true: The higher the probability, the better the explanation, and so the more likely it is to be true.

The best possible causal explanation then is one that inter alia raises the causal probability as far as it can go, namely to 1. That is, it is a deterministic causal explanation conforming to the deductive–nomological model. And that is my final reason for requiring causes to raise the probability of their effects. It explains, what is otherwise a complete mystery, why we would obviously always prefer a deterministic explanation (other things being equal) if we could get it—so obviously that until the pioneering work of Salmon and others, indeterministic causal explanation was, as he says, thought to be a contradiction in terms. But in showing that it isn't, it is important not to get carried away and weaken the concept of causation so much that the appeal of deterministic causal explanation is lost sight of, the principle of inference to the best explanation mysteriously separated from the maximum likelihood principle, and the rationale destroyed of acting for the sake of the action's effects. That cannot be right, and it is in no way necessary to a probabilistic theory of causal explanation.

REFERENCES

Jeffrey, R. C., 1980, How is it reasonable to base preferences on estimates of chance?, in Mellor, D. H. (ed.), *Science, Belief and Behaviour*, Cambridge, Cambridge University Press, 179–187.

Mellor, D. H., 1976, Probable explanation, *Australasian Journal of Philosophy* 54, 231–41.

Salmon, W. C., Jeffrey, R., and Greeno, J., 1971, *Statistical Explanation and Statistical Relevance*, Pittsburgh, University of Pittsburgh Press.