THE JOURNAL OF PHILOSOPHY

VOLUME LXXXII, NO. 2, FEBRUARY 1985

PROBABILITIES AND CAUSES

THIS paper is about the connection between probabilities and causal explanation. I shall start by describing two familiar views on this matter. I shall call them the *standard view* and the *statistical-relevance view*. After that I shall introduce a different way of understanding the connection, which I shall call the *evidential view*. I shall conclude by arguing that once the evidential view is taken into account, the standard view and the statistical-relevance view lose their attractions.

I. THE STANDARD VIEW

On the standard view, the connection between probabilities and causal explanation is simple: one event explains another just in case the former gives the latter a high probability.[1] The higher the probability, the better the explanation. So the ideal case is the limiting one, where the explanans event determines the explanandum, and so allows us to deduce its occurrence; but one still has explanation, albeit of a weaker variety, when the probability falls short of unity, provided it is still reasonably high.

Thus, to take an example, if *all* people with streptococcal infections who take a certain dose of penicillin recover, one will be able to give a deductive explanation of why a particular sufferer who takes the dose recovers. But even if only 90% recover one can still explain, albeit slightly less satisfactorily, someone's getting better in terms of their taking the penicillin.

This standard view[2] faces a well-known difficulty. Suppose, as

[1] A terminological point. It will simplify the exposition if I can assume that all explanation of particular events is causal, that explaining one event by another is always a matter of showing the first to be caused by the second. As it happens, I think this assumption is right. But I shall not be concerned to defend it as such, so I would merely ask those who disagree to accept it as a terminological convenience, and remember that when I say 'explanation' I should be read as talking specifically about *causal* explanation.

[2] The "standard view" is of course the view developed by Carl G. Hempel. See his *Aspects of Scientific Explanation* (Glencoe, Ill.: The Free Press, 1965), ch. 12, pp. 376–411. However, it will simplify matters without, I hope, begging questions to work with an ideal type and avoid the details of Hempel's account.

0022-362X/85/8202/0057$01.80

above, that 90% of streptococcally infected people recover (R) upon taking penicillin (P): Prob(R/P) = 0.9. But suppose also that in a specific subclass of sufferers—those more than 70 years old, say—this proportion is radically lower. Suppose, for instance, that the proportion of recoverers among old (O) people who take the drug is only 10%: Prob($R/P\&O$) = 0.1.

And now consider the case of an old person who takes the penicillin and *does* recover. We've still got the probabilistic generalization applying to penicillin takers in general—Prob(R/P) = 0.9—and it covers this old person—he or she took penicillin (P). So, for all that's been said so far, it seems that we ought to be able to explain this particular recovery by the ingestion of the drug.

But surely we don't want this. Suppose that the same old person *hadn't* recovered, as was in fact very likely. Then we could have explained that too. For we have another generalization—Prob ($-R/P\&O$) = 0.9—and initial conditions—$P\&O$—in terms of which the *non*recovery could then have been given high probability.

What seems wrong here is the possibility of our explaining some outcome (the recovery) in the light of certain features of the actual situation, when we could have explained the opposite outcome (the nonrecovery) just as well if that had occurred, by mentioning other features of the situation.

But there is a natural enough way of ruling out this possibility. We can impose a requirement of *maximal specificity*: we can require that initial conditions mentioned in our explanation include *all* the relevant features of the situation. This then will stop the explanandum getting a high probability simply because certain relevant factors (the person is old) have been omitted from the explanans.

This maximal-specificity requirement is equivalent to the demand that the probabilities entering into our explanation should derive from a "homogeneous" partition of our reference class. Let me explain. In our example the original reference class was the class of all infected people. We then "partitioned" this reference class into the subclasses of those who took penicillin and those who didn't. But this partition failed to be "homogeneous" for recovery, in the sense that the probability of recovery depended not just on whether one took penicillin, but also on whether one was old. So we needed to partition again, into subclasses distinguished by age. A homogeneous partition for some outcome is a partition such that no further subdivisions will further alter the proportions with which that outcome is found within subclasses. Penicillin and age will give us a homogeneous partition for recovery if there are no further factors on which the probability of recovery depends.

The advantage of putting the maximal-specificity requirement in terms of homogeneity is that this makes explicit the precise sense in which the standard view should require an explanation to include all *relevant* factors. In our example, it wouldn't matter if the explanation omitted factors, like hair color, say, which were true of the person in question, but made no difference to the probability of recovery. What needs to be ensured is specifically that we include all those factors which *do* make such a probabilistic difference—that is, which are required to yield a homogeneous partition of the original reference class.

II. THE STATISTICAL-RELEVANCE VIEW

An alternative to the standard view is the view that the explanation of one event by another does not require that the former give the latter a *high* probability, but simply that it make it *more* probable than it would otherwise have been.[3]

On this view (the *statistical-relevance*[4] view henceforth), for someone's cigarette smoking (S) to explain their lung cancer (C) it is not required that smoking make cancer highly probable: Prob $(C/S) >> 0.5$, but only that it increase the probability: Prob(C/S) > Prob$(C/-S)$. It doesn't impugn the explanatory connection, on this view, that even among smokers the probability of cancer remains well below 50%: provided smoking increases the probability—from 1% to 10%, say—we have enough to allow explanations of cancer in terms of smoking.

As with the standard approach, there is an obvious (and indeed not unrelated) objection to the statistical-relevance approach. According to the statistical-relevance approach, the requirement for one event (A) to cause another (B) is that Prob(B/A) > Prob$(B/-A)$. Equivalently, that Prob(B/A) > Prob(B); or, again, that Prob$(A\&B)$ > Prob(A)Prob(B). In short, the requirement is that A and B be "correlated," or "statistically associated," with each other. How-

[3] This view is distilled from the writings of Patrick Suppes and Wesley Salmon. See Suppes, *A Probabilistic Theory of Causality* (Amsterdam: North-Holland, 1970), and Salmon *et al.*, *Statistical Explanation and Statistical Relevance* (Pittsburgh, Pa.: University Press, 1971). But again it will simplify matters to work with an ideal type. Indeed I should make clear that my "statistical-relevance view" differs from both Suppes's and Salmon's accounts on a number of points of substance.

[4] The term 'statistical-relevance' might suggest that the important difference between this approach and Hempel's is that it excludes irrelevant explanatory factors whereas Hempel fails to do so. But this is misleading. If Hempel's explicit formulations do sometimes leave room for irrelevant factors, this is a trivial matter, and easily remedied. One simply adds to the maximal-specificity requirement: include all relevant factors, the obvious corollary: omit irrelevant factors. The real difference between Hempel and the "statistical-relevance" alternative has nothing to with relevance as such; it is simply the difference between high probabilities and increased probabilities. Cf. sec. 5 of Salmon, "Statistical Explanation," in Salmon *et al.*, *op. cit.*

ever, as generations of social science students know, "Correlation doesn't prove causation."

Thus, to take an example, although there is certainly a well-established correlation between smoking and cancer, this doesn't necessarily mean (as the cigarette companies wistfully point out) that smoking *causes* cancer. It may yet be that there is some hidden factor (a genetic predisposition, say) which, on the one hand, causes cancer, but which also, as an independent side effect, induces people to smoke. This would give us a perfectly good correlation between smoking and cancer, even though the first didn't cause the second.

To deal with this difficulty the statistical-relevance account needs to introduce a "no screener-off" requirement. Suppose we have an association between C (cancer) and S (smoking). That is, $\text{Prob}(C/S) > \text{Prob}(C)$. Then we shall say that some factor G (think of the genetic predisposition) "screens off" the association between C and S, if $\text{Prob}(C/S\&G) = \text{Prob}(C/G)$, and $\text{Prob}(C/S\&{-}G) = \text{Prob}(C/{-}G)$. That is, the genetic predisposition would screen off the association between smoking and cancer if, among people who had the genetic predisposition, cancer were no more frequent in the smokers than in the nonsmokers and if the same were true among those who lacked the genetic predisposition.

If some G did so screen off cancer from smoking, then clearly the cigarette companies would be right, and the correlation between smoking and cancer would merely be due to the fact that smoking was more frequent among those who were going to get cancer anyway. Conversely, if there isn't any such G to explain away the original association then it seems to follow that there really is a direct causal link between the two. So let us take the statistical-relevance account to hold that A causes B just in case $\text{Prob}(B/A) > \text{Prob}(B)$ and there is no C which screens off B from A.

Why didn't the "standard view" need a "no screener-off" requirement? That is, why wasn't the standard view open to the objection that the probabilities appearing in its explanations might be "spurious" ones which made mere symptoms appear as causes? It is true that the standard view required high probabilities where the statistical-relevance view required only increases. But high probabilities can be spurious just as much as low ones.

The answer is that the standard view has already ruled out screeners-off by imposing the homogeneity requirement. The homogeneity requirement says that no further factors should make a difference to the probability of B. And so *a fortiori* it rules out any factors that will, so to speak, change $\text{Prob}(B/A)$ from being bigger than $\text{Prob}(B/{-}A)$ to being equal to it.

But it is worth noting, and it will be important for what follows, that the converse implication does not hold. Although the homogeneity requirement implies the no-screeners-off requirement, the no-screeners-off requirement does not imply homogeneity. The reason is that the no-screeners-off requirement rules out only a quite specific kind of probability-altering factor, namely, those factors that, in altering the probability of *B*, make the association between *B* and *A* "disappear." So the nonexistence of screeners-off still leaves it open that there may be other probability-altering factors in the offing and that the homogeneity requirement is therefore not satisfied.

Thus, to return to our original example, although being old altered the probability of recovery and so stopped the simple penicillin/no-penicillin partition from being homogeneous for recovery, it doesn't follow that it screens recovery off from penicillin. It's still perfectly possible (and indeed no doubt true) that among the old, as among the young, recovery is more likely for those who take penicillin. And consequently the fact that being old alters the probability of recovery does not, on the statistical-relevance account, discredit the suggestion that the penicillin caused the recovery, when somebody who takes penicillin gets better.

Note that this now gives us something of a test case. I took it, in the first section, that when an old person who took penicillin recovered, we wouldn't be able to explain this. After all, what actually happened was terribly unlikely, given a full specification of the circumstances. But the statistical-relevance account disagrees. Provided that being old doesn't actually screen off the recovery from the penicillin, provided that the chance of recovery is still increased by taking penicillin, then, even if the probability of recovery remains very low, when we do get a recovery, we should deem it to be caused by the penicillin.

This makes it clear why the standard view requires homogeneity, where the statistical-relevance view is happy to settle for no screeners-off. The standard view says we have an explanation if all the factors present in the circumstances give the explanandum a high probability. Obviously all relevant factors need to be taken into account to determine whether this is so. But the statistical-relevance account doesn't put any demands on the absolute value of the explanadum's probability. All it requires is that this probability be increased by the presence of the putative cause, and that this increase not be screened off. And so the statistical-relevance account can safely ignore the question of whether any other factors have a similar influence.

However, though it is clear enough what is in dispute between

the standard and statistical-relevance views, it is not at all clear which view is right. Indeed the situation seems downright puzzling. Even in the test case, *both* views seem highly plausible. Intuitions pull strongly in both directions. Surely unlikely events aren't causally explicable. But isn't an increase in probability *per se* enough to indicate causation? This tension should make us suspect that there is more to the dispute than meets the eye, that underneath the surface of the discussion a number of different issues are present. So indeed I shall shortly argue. But in order to do this I first need to elaborate a yet further view on the connection between probabilities and causes.

III. THE EVIDENTIAL VIEW

So far I have been implicitly presupposing that, whatever the exact form of the relationship between probabilities and causes, causal relationships are somehow *constituted* by probabilistic relationships. That is, I have intended that both the standard and the statistical-relevance views be understood to be claiming that one event's causing another just *is* a matter of there being certain probabilistic generalizations relating events of those types, analogously to the way that on more old-fashioned views causation is a matter of there being certain deterministic generalizations relating the event types. Henceforth I shall call views of this kind *definitional* views. I want to contrast definitional views with an *evidential* approach to the relationship between probabilities and causes.

(Perhaps it is implausible that anybody nowadays might suppose that causation can be *defined* in terms of probabilities, or indeed in terms of anything else. But, as will become clear shortly, it's not really the possibility of a full analysis of causation that's at issue, so much as opposing views on the connection between probabilities and the substance of causation.)

To illustrate the evidential view, let us pretend for a while (indeed for the rest of this section) that the "old-fashioned" deterministic view of causation referred to above is in fact correct: let us pretend that causation is necessarily deterministic. I want to show how, even given this assumption, probabilities (that is, probabilities other than 0 and 1) can tell us about causes. The idea will be that such probabilities can still be seen as *evidence* for the underlying presence of such deterministic relationships. (The assumption that causation is deterministic is not essential to the evidential approach. I return to this in the next section.)

It will be convenient to adopt John Stuart Mill's terminology for analyzing deterministic causation, and call some factor A a *cause* of some B just in case it is one of a set of conditions that are jointly

and minimally sufficient for B. In such a case we can write $A\&X \rightarrow B$. In general there will also be other sets of conditions minimally sufficient for B. Suppose we write their disjunction as Y.

If now we suppose further that B is always determined when it occurs, that it never occurs unless one of these sufficients sets (let's call them B's *full causes*) occurs first[5], then we have

$$A\&X \vee Y \leftrightarrow B$$

Given this equivalence, it is not difficult to see why A's causing B should be related to A's being correlated with B. If A is indeed a cause of B, then there is a natural inference to $\text{Prob}(B/A) > \text{Prob}(B/-A)$: for, given A, one will have B if either X *or* Y occurs, whereas *without* A *one will get* B *only with* Y. *And conversely it seems that* if we do find that $\text{Prob}(B/A) > \text{Prob}(B/-A)$, then we can conclude that A is a cause of B: for if A didn't appear in the disjunction of full causes which are necessary and sufficient for B, then it wouldn't affect the chance of B occurring. This then gives the basic idea of the evidential view. But of course it is a bit too quick. To take the latter inference first, A's being a cause of B isn't the only way in which it can get to be correlated with it; for even if A is not itself a member of a full cause, it may be statistically associated with factors that are, and this in itself will give rise to a ("spurious") correlation between A and B. And conversely, if not so obviously, even if A *is* a cause of B it may be (though it would be something of a freak) negatively associated with elements of other full causes of B to just the extent required to leave us with a null correlation between A and B.

In effect what we are worried about here is the existence of factors that might screen off the correlation (or lack thereof) between A and B. But the evidential view can cope with this possibility by taking into account a more complicated set of statistical relationships. Suppose $[D]$ is the set of all the other causes of B which are

[5] It is worth being clear about what is involved in assuming that some B is in this sense determined. We are assuming that there is, so to speak, determinism *in between* $A\&X \vee Y$ and B. This does not commit us to the view that the whole universe is deterministic. This point is important for the status of such probabilities as $\text{Prob}(B)$, $\text{Prob}(B/A)$, etc. I shall want to continue talking about such probabilities as general facts involving event types, and I shall want it to be possible for these probabilities to be other than 0 or 1 even when B is determined on each occasion of its occurrence. It is arguable, however, that serious probabilities other than 0 and 1 need to be underpinned by real indeterministic chances. But I don't need to take issue with this: we can still get serious nondegenerate probabilities even if B is determined by A, X, Y, etc., provided only that the occurrence of A, or X, or Y, etc., is itself the outcome of some real chance process, or that they are in turn determined by factors whose occurrence is the outcome of some chance process . . .

associated with A. What we need to do is partition the original reference class by all combinations of presence and absence of the factors in $[D]$, and then consider whether A is correlated with B *within* any of the cells of this partition. If A is so correlated then it must now be a cause, because there aren't any "remaining" associations within any cell left to produce "spurious" correlations between A and B. And, conversely, if there aren't any correlations between A and B left within cells, if any original correlation between A and B has been screened off by some member of $[D]$, then it follows that A really isn't a cause after all.

For those who would like to look more closely at these inferences there is an appendix that spells out the steps involved. Perhaps one brief observation will be helpful at this point. Partitioning by all *associated* causes $[D]$ is not the same as partitioning by all other causes *tout court*. If the latter was required then, given the determinism we are assuming, we wouldn't be in a position to make inferences about causation until we had a partition that gave B a probability of either 1 or 0 within every cell. But since we don't have to control for all other causes—since, that is, taking all possible screeners-off into account doesn't require that we have a fully homogeneous partition—it *is* possible to make cogent causal inferences from probabilities other than 0 or 1.

IV. ADJUDICATION

a. *Deterministic Situations.* So now we have three different approaches: standard, statistical-relevance, and evidential. Which is right?

On the surface, perhaps, the statistical-relevance and the evidential accounts look similar. They both focus on increased probabilities rather than high ones. But in intent they are quite distinct. We can best bring this out by considering once more the kind of deterministic situation assumed in the last section. Thus let us suppose, again for the sake of the argument, that lung cancer (C) is always determined when it occurs, either by smoking (S) in conjunction with certain (presumably metabolic) preconditions X or by certain other conjunctions of factors Y which don't include smoking. It might seem that there will be a stand-off between the statistical-relevance and the evidential approaches in such cases. Both sides will agree about the structure of probabilistic and deterministic generalizations relating C, S, X, and Y. Both will agree that S causes C. But where the statistical-relevance approach will say that this is because S is correlated with C and this correlation isn't screened off by anything else, the evidential approach will say that this is because S is part of a full deterministic cause of C (and

that the statistical relationships, including the absence of screeners-off, are merely indirect evidence for this).

It doesn't look as if there is much in it. But the statistical-relevance approach is in fact in trouble here. Consider an actual person who smokes and then gets cancer. According to the statistical-relevance approach, the smoking caused the cancer. But what if that person wasn't *X*, but was *Y*? It seems clear that in that case (let's call it the "case of the misleading cigarettes") the smoking wouldn't have caused the cancer—that though the person smoked, that wasn't why he or she got cancer.

And this of course is what the evidential approach will say. The probabilities are only indirect evidence for the underlying causal connection, and so smoking and cancer can be probabilistically connected in general without being causally connected in the particular case.

But the statistical-relevance approach seems to lack the room to say this. If causal connections are *constituted* by the relevant probabilistic associations, then it seems to follow from the general connection between smoking and cancer that the former causes the latter in the misleading particular case.

This case of the misleading cigarettes brings out an ambiguity that I have been slurring over so far. One can read '*A* causes *B*' as a relationship between types: 'being *A* is the kind of thing that causes *B*', 'on some occasions *A* gives rise to *B*', '*A can* cause *B*'. But one can also read it as a relationship between tokens: 'on this occasion *A*'s being instantiated caused *B*'s being instantiated'. In many contexts the distinction is unimportant—the token relationships between instantiations of *A* and *B* will obtain just in case the type relationships between the kinds as such obtain.

But for the evidential approach the distinction does matter: the type relationship that it takes to be evidenced by the probabilities (that *S* is a member of a set of conditions sufficient for *C*) doesn't guarantee that the relevant token relationship will obtain whenever a particular *S* is followed by a particular *C* (as when some other members of the sufficient set in question are absent). Which of course is why the evidential approach can deal with the misleading-cigarettes case.

But a corresponding move does not seem open to defenders on the statistical-relevance approach. By viewing the probabilities constitutionally rather than evidentially, it seems that they will be forced to conclude that a causal relationship between tokens obtains whenever we have tokens of kinds related by the appropriate probabilities.

Perhaps defenders of the statistical-relevance approach will feel inclined at this point to appeal to a distinction between type and token *probability* relationships. Of course, they will agree, in a general (type) sense smoking increases the probability of cancer. But it doesn't follow that it increases the probability in the particular case at issue, since, after all, the probability in the particular case was already 1, because of the presence of Y, quite independent of the smoking. Thus if we tie causation to token probabilities rather than type probabilities, we can avoid saying that the person with the misleading cigarettes got cancer because of smoking.

But note that this ploy means that the statistical-relevance approach has abandoned the distinctive no-screeners-off requirement. The switch from type probabilities to token probabilities, to probabilities that reflect the real chance of the outcome given all the circumstances to hand in the particular case, demands that our probabilities come from *homogeneous* reference classes, not just from classes that take into account all screeners-off.

And this then means that, in the deterministic case, the statistical-relevance approach collapses into the standard one, with both demanding that satisfactory causal explanations take all relevant factors into account. Indeed in the deterministic case both then leave us with the outcome always having a probability of either 0 or 1. And so neither the statistical-relevance nor the standard account has anything to say about deterministic situations which was not already said by the orthodox deductive covering-law model.[6]

b. *Indeterministic Situations.* Let me turn now to indeterministic situations. Some readers will no doubt feel this is long overdue. Perhaps the evidential view is the view demanded by deterministic assumptions. But the world isn't in fact deterministic. And so isn't the important quesiton whether the standard or the statistical-relevance view deals best with indeterminism?

But this would be too quick. The evidential view has things to say about indeterministic situations too.

This might seem a bit puzzling. In order to explain it we need to do a bit of unraveling. Leave the evidential view out of it for a moment. I have been speaking for the most part as if there were

[6] Couldn't the statistical-relevance approach (or, for that matter, the standard approach) maintain that we could have explanations underpinned by nondegenerate probabilities in what were actually deterministic situations as long as we believed (with good reason?) that our probabilities were homogeneous? But though we would then "have explanations" insofar as we believed we could explain, there is a clear sense in which they wouldn't necessarily *be* explanations—for they would be downright mistaken whenever we had cases like that of the misleading cigarettes.

two definitional views of causation, the standard and statistical-relevance views. But we have also at various points come across what is in effect a third definitional view, namely, the "old-fashioned" deterministic view, according to which causation necessarily requires a deterministic generalization linking the event types involved. So from now on let us take it that there are three definitional views in the field: two indeterministic (the standard and statistical-relevance) and one deterministic ("the old-fashioned").

The evidential view is not in competition with definitional views. It is a view not about the constitution of causation, but about the possibility of a certain kind of indirect evidence for causal relationships. In section III, it is true, I introduced the evidential view by putting it in harness with the old-fashioned deterministic view of the constitution of causation. But, as I said at the time, this link was not essential. What I now want to show is that one might still need to adopt an evidential attitude even if one is taking an indeterministic view of the constitution of causation.

To see this, let us temporarily assume, for the sake of argument, such an indeterministic view of the constitution of causation. And let us agree that cancer, say, is not a deterministic outcome of prior circumstances. Even then, I want to argue, one needs to recognize the possibility of taking an evidential attitude to the connection between probabilistic generalizations and causal conclusions.

For it certainly seems intuitively clear that, even with these assumptions, one will want to allow that probabilities can tell us about causes even though the reference classes from which we are getting the probabilities are not homogeneous—surely we want to allow that we can have good grounds for thinking that smoking causes cancer even though the probabilities we take to show this would be altered by the introduction of further factors. Of course we need to recognize that here, as always, a sound inference requires that we take note of *some* of the further factors relevant to getting cancer. For we need to take into account possible screeners-off (which would include such factors as genetic make-up, psychological stress, diet, etc.). But this isn't the same as needing to take into account *all* relevant factors. No doubt there are as yet unknown metabolic differences between people which affect the chance of (without necessarily determining) cancer. But provided that we have reason to believe that these differences aren't themselves correlated with smoking, as they needn't be, and so can't screen off cancer from smoking, our ignorance of what they are and what exact effect they have on the probabilities clearly doesn't invalidate the inference to a causal conclusion.

The importance of this point, that even in the indeterministic situation inferences to causes demand no screening off rather than homogeneity, is that it shows that the evidential view is still at work. Consider once more the case of the misleading cigarettes. As originally described, this presupposed determinism. But we can get an indeterministic analogue. Let smoking plus X, and the various factors in Y, give a homogeneous partition for cancer, but now one that doesn't always determine cancer or its absence, but merely makes it probable to various degrees. Then one might have a particular person who smoked, and had a raised chance of cancer, and got cancer—but not because of the smoking, for they lacked X, but because of Y. Although smoking raises the probability of cancer, and this association isn't screened off, it doesn't follow that when a smoker gets cancer the smoking caused the cancer.

This shows that if you are concerned with the realization of causation in particular cases, with information that will tell how the causes are constituted, then, as in the deterministic situation, you are going to need homogeneous reference classes. It is only when you have such homogeneous probabilities that you can be sure that the factors you allude to by way of explanation in the particular case are those which in *that* case made the result more probable, not just factors that have the appropriate general (type) relationship with the outcome.

And, conversely, this shows that if you are happy with merely unscreened-off probabilities, rather than homogeneous ones, then you must be adopting an evidential rather than a definitional view of the relationship between probabilities and causes. The probabilities can't then be the substance of the causal relationships, but merely an indication of underlying connections.

It might still be felt that I am avoiding the main issue. Perhaps definitional views of causation require homogeneous probabilities across the board, in indeterministic as well as deterministic situations, and perhaps to that extent the statistical-relevance view should never have mentioned the weaker no-screeners-off requirement in the first place. Perhaps, that is, the no-screeners-off idea is of importance only to the evidential view. But doesn't there still remain the original question: what is the right definitional account of causation for indeterministic situations? That is, even if indeterministic accounts—the standard and statistical-relevance accounts alike—require homogeneous probabilities, isn't there still the question of whether these probabilities should be high probabilities or merely increased ones?

There is, however, a good reason why I have spent so long on the

evidential view before turning to this question, namely, that once we have distinguished the evidential view, and seen that it can apply in both deterministic and indeterministic situations, there is far less reason than before to suppose that *either* the statistical-relevance view *or* the standard view is correct.

Let me deal first with the statistical-relevance view. What I want to suggest is that the statistical-relevance view seems plausible in the first place only because it gets confused with the evidential view.

The statistical-relevance view implies that somebody who has only a 10% chance of getting cancer, given *all* the relevant factors present in the particular situation, can be caused to get cancer by those factors. I find this counterintuitive. Just this person might well (nine times in ten) not have got cancer. Surely then we can't explain why he or she did?

This was the intuition I appealed to in section I when I first argued that the maximal-specificity requirement needed to be added to the standard view. But this intuition is rather unfashionable nowadays. And of course the reason is the further thought, familiar from the writings of statistical-relevance theorists, that, if smoking increases the probability of cancer, even if only from 1% to 10%, then this in itself indicates that smoking causes cancer.

But I think we should resist fashion and stand by the original intuition. For I think fashion has been misled by a failure to distinguish the statistical-relevance from the evidential view, and has therefore misread the significance of increased probabilities. It is right to hold, with the evidential view, that increased probabilities are important. But it is wrong to conclude, with the statistical-relevance view, that they can constitute causation.

According to the evidential view, increased probabilities are exactly what we need in order to *infer* the existence of underlying causal connections. That smoking increases the probability of cancer from 1% to 10% indeed gives us good reason to conclude that smoking causes cancer. But on the evidential view we don't need to suppose that the causation itself is *constituted* by a mere 10% chance. The conclusion of our inference, on this reading, could be that there existed certain other unknown factors which, when they occurred together with smoking, gave cancer a high probability.

Thus the evidential view allows us both to stand by the original intuition that causation requires high probabilities and to accommodate the insight that increased probabilities are important. And it seems to me that, once this is recognized, the statistical-relevance view loses its plausibility. Once we can give a nondefinitional account of the feeling that increased probabilities are important,

there remains little to be said in favor of the statistical-relevance view.

What now of the standard view? Here things are less clear. It might seem that if the statistical-relevance view is out of the running, then the standard view must be the right account of indeterministic causation. However, we shouldn't take it for granted in the present context of argument that there is any "right account" of indeterministic causation. Perhaps, that is, all indeterministic accounts of causation should be rejected, in favor of the old-fashioned deterministic account.

It is important to recognize here that the question of whether *causation* is definitionally indeterministic or not is in principle quite independent of whether the *world* is indeterministic or not. There is every reason to believe that the world is indeterministic. But one can allow this, and still stand by a deterministic view of causation. All that would follow would be that, to the extent that things aren't determined, they aren't caused either.[7]

What arguments are there for allowing indeterministic causation? I suspect that many people have been swayed by the following line of thought. "It is scarcely deniable that probabilities have *something* to do with causes. A huge amount of empirical research, by medical investigators and agricultural scientists, not to mention psychologists, economists, and sociologists, proceeds from probabilistic data to causal conclusions. Clearly this isn't all just based on a conceptual mistake. So surely we should accept that causes can be constituted by probabilities as well as by deterministic certainties." But the evidential view shows that this reason for rejecting the old-fashioned deterministic view of causation is mistaken. For the evidential view, as originally illustrated in section III above, shows how it would still make perfect sense to infer from probabilities to causes even if one assumed that causes were always constituted by underlying deterministic connections.

But this scarcely decides the matter. Perhaps there are independent reasons, apart from the cogency of basing causal conclusions on probabilistic data, for thinking causation needn't be deterministic. And on reflection it does perhaps seem implausible to say that indeterminism *per se* precludes causation. Suppose that roulette

[7] One could even, though there wouldn't be much point, do it the other way round: one could combine an acceptance of an indeterministic definitional view of causation with the thesis that the world is in fact deterministic, holding then that even if the world had been indeterministic, such indeterminism wouldn't have implied an absence of causation.

wheels are really indeterministic (suppose, that is, that quantum effects always have an influence on where the ball lands). Should that stop us saying (on most evenings) that the house is caused to win by the presence of a zero on the wheel? (Which, note, gives it an extremely high probability of winning, not just a higher probability than it would otherwise have.)

But neither is this reflection conclusive on the other side. The old-fashioned view doesn't need to deny that there is causation in the casino. If roulette wheels are really indeterministic, then the physical nature of the set-up determines that there is a certain single-case probability (a *chance*) of the house winning. And so the old-fashioned view can say that what is caused in the casino is the very high chance of the house winning over the evening. This causation is itself deterministic. And that, so the old-fashioned thought would go, is all the causation there is. Once the high chance has been explained, there is no further explanation of why the house won—after all, given just the same chance on another night, the house might not have won, so why suppose we can explain its winning, as opposed to its being very likely that it would?

It might seem that there is little to choose at this stage between the standard and old-fashioned ways of admitting causation into the casino. However there is a strong reason, given what has gone before, for preferring the old-fashioned account. This relates once more to the possibility of drawing causal conclusions from inhomogeneous but screener-off-free probabilistic data.

Note that once we relax deterministic assumptions, we cannot continue to assume that an unscreened-off $\text{Prob}(B/A) > \text{Prob}(B)$ shows that A is (part of) a determining condition for B. What it shows (as the appendix explains) is rather that A is part of a homogeneous reference partition for B. But—and this is the crucial point here—this doesn't mean that A is (part of) a set of factors that makes the chance of B *high*. And so, if the standard view were right, and causation required that A (plus some X) made the chance of B high, then there would be no valid inference from an unscreened-off $\text{Prob}(B/A) > \text{Prob}(B)$ to the conclusion that A is a cause of B.

There are two ways to go here. (1) We could reject the assumption that unscreened-off correlations indicate causal conclusions; or (2) we could insist on the old-fashioned view in place of the standard view, and maintain that in indeterministic situations it is only the chances of undetermined events that get caused, rather than the events themselves. I take it that (1) is unacceptable—intuition, in-

stantiated in a whole tradition of empirical research, tells us that *something* about causal conclusions must follow from un-screened-off correlations. So I maintain (2), saying that when we have an un-screened-off $\text{Prob}(B/A) > \text{Prob}(B)$ in an indeterministic situation, what we can infer is that A is (part of) a deterministic cause of the chance of B.[8]

This rejection of the standard view for the old-fashioned view might seem a somewhat unsatisfactory maneuver forced on me only because of my rejection of the statistical-relevance view. For note that the statistical-relevance view, unlike the standard view, does allow us to draw causal conclusions from un-screened-off correlations in indeterministic situations: increased probabilities, by contrast with high ones, are indeed reliably indicated by un-screened-off correlations.

Even so, I think it is right to reject the statistical-relevance view. As I have argued, the idea that mere increase in probability constitutes causation is in itself counterintuitive; and what plausibility there is to the statistical-relevance view can be explained away as due to its conflation with the evidential view. We now see that the alternative to the statistical-relevance view is to construe causation in indeterministic situations as the deterministic fixing of chances by prior circumstances. Even if this seems awkward at first sight, I don't see that there is anything definite to be said against it. And so I see no reason not to stand by the rejection of the statistical-relevance view, and to return to the old-fashioned deterministic view of causation across the board.

DAVID PAPINEAU

University of Cambridge

APPENDIX

Let me start by going back to the assumptions of Section III. That is, let me take it not only that causes are deterministic, but also that B is determined whenever it occurs (so we can write: $A\&X \vee Y \leftrightarrow B$). I want to show that, given these assumptions, A will be a cause of B if and only if there is an un-screened-off probabilistic association between A and B.

This argument has two stages. I shall first consider the simple case where A is probabilistically independent of X and Y and truth functions thereof. Then I shall proceed to the more complicated case where such independence is not assumed.

[8] I would like to thank Crispin Wright for help with the argument of the last two paragraphs.

The simple case. Note that

$$\begin{aligned}\text{Prob}(B/A) &= \text{Prob}(A\&X \vee Y/A)\\ &= \text{Prob}(A\&X \vee A\&Y)/\text{Prob}(A)\\ &= [\text{Prob}(A\&X) + \text{Prob}(A\&Y) - \text{Prob}(A\&X\&Y)]/\text{Prob}(A)\\ &= \text{Prob}(X) + \text{Prob}(Y) - \text{Prob}(X\&Y)\end{aligned}$$

(since A is independent of X, Y, and $X\&Y$)

$$= \text{Prob}(X \vee Y)$$

So if $\text{Prob}(B/A) > \text{Prob}(B)$, then $\text{Prob}(X \vee Y) > \text{Prob}(B)$. That is, $\text{Prob}(X \vee Y) > \text{Prob}(A\&X \vee Y)$. Which means that there is indeed some X that in conjunction with A figures in the list of B's full causes. [If there weren't—if, so to speak, that X were the impossible condition—then $\text{Prob}(X \vee Y)$ would have to equal $\text{Prob}(A\&X \vee Y)$.] Thus $\text{Prob}(B/A) > \text{Prob}(B)$ implies that A is a cause of B.

Conversely, if A is a cause of B—if there is a "non-impossible" X such that $A\&X$ figures in the list of B's full causes—then $\text{Prob}(B/A)$, which is equal to $\text{Prob}(X \vee Y)$, must be bigger than $\text{Prob}(B)$, which equals $\text{Prob}(A \,\&\, X \vee Y)$. The only way we could have an equality would be if $X \vee Y$ never occurred in the absence of A, but this is ruled out by the independence assumption. [A slight complication. If A *always* occurred, then $X \vee Y$ would never occur in its absence, even if $X \vee Y$ was probabilistically independent of A, and we would have $\text{Prob}(B/A) = \text{Prob}(B)$. Perhaps we should say that something that "always" occurs, whatever precise modal significance we attach to this, cannot qualify as a cause. But what then about something that "almost" always occurs, in the measure-theoretic sense? Again, this will give $\text{Prob}(B/A) = \text{Prob}(B)$. Here I think we should simply relax the evidential thesis, and admit that to this extent causes won't automatically show up in probabilities.]

I turn now to the *second stage* of the argument, where we relax the independence assumption that A is independent of X and Y and truth functions thereof. Now we need to "control" for the elements in X or Y (remember that X and Y will themselves generally be disjunctions or conjunctions of factors) which are associated with A and which therefore might be "confounding" the probabilistic relationship between A and B. That is, we need to look separately at classes of cases similar with respect to the presence or absence of such confounding elements, and see whether A is still probabilistically associated with B within such classes. The claim, then, is that A is a cause of B if and only if such "un-screenable-off" associations remain.

The strategy here is simple enough. The original reference class is partitioned by means of combinations of presence and absence of factors from X and Y, in such a way that A is independent of the "remaining" parts of X and Y and truth functions thereof within each cell of the partition. Then in effect we repeat the original argument within each cell of the partition. Given the assurance that within each cell A is independent of the "remain-

ing" causes, we can conclude that (a) if there are some cells in which A is still probabilistically associated with B, then it must be a cause (there must be some condition together with which A appears on the list of B's full causes), and (b) if there are no cells in which A remains probabilistically associated with B, then it can't be a cause (modulo the qualification made earlier about As which "almost" always occur, now relativized to cells). If A is associated with B within some cells but not within others, then A and the factors defining those cells are said to "interact" in causing B. This will occur if A is associated with (and we therefore have to "control" for) some factor which is a necessary part of X.

All the above argument is set within the deterministic assumptions of Section III. To generalize the explication of the evidential approach to indeterministic situations, we simply need to "replace" all undetermined events with their chances of occurrence, and then proceed as before. That is, we can regard the aim of the evidential approach in indeterministic situations to be the identification of the *deterministic* laws relating the chances of undetermined events to antecedent conditions. The presence of an un-screened-off correlation between A and B now indicates that A is one of the conditions that plays a part in determining the chance of B. Note that this does not mean that A plays a part in determining that B is *highly* probable, for the relevant chance of B need not be high.

DP

TRANSLUCENT BELIEF*

ONE of the more perplexing problems in the philosophy of language involves the interpretation of belief ascriptions—sentences of the form:

R believes that p.

where the expression that replaces p (the content of the ascription) has the syntactical form of a sentence. The problem is to say how such contents are to be interpreted.

Philosophical lore has it that our alternatives are exactly two: either an ascription of belief is transparent, or it is opaque.[1] If it is transparent, the content is just the sentence it appears to be, and its interpretation is straightforward. If it is opaque, this is not the

*I am grateful to Israel Scheffler, Nelson Goodman, and Jonathan Adler for helpful comments on an earlier draft of this paper, and to the University Research Council of the University of North Carolina at Chapel Hill for financial support.

[1] Some may be one, some the other. Indeed, some tokens of a single type may be transparent, others opaque. But tradition has it that unless a token is transparent, it is opaque. There is no third alternative.

0022-362X/85/8202/0074$01.70