CARL G. HEMPEL

Inductive-Statistical Explanation

1 | Inductive-Statistical Explanation

As an explanation of why patient John Jones recovered from a streptococcus infection, we might be told that Jones had been given penicillin. But if we try to amplify this explanatory claim by indicating a general connection between penicillin treatment and the subsiding of a streptococcus infection we cannot justifiably invoke a general law to the effect that in all cases of such infection, administration of penicillin will lead to recovery. What can be asserted, and what surely is taken for granted here, is only that penicillin will effect a cure in a high percentage of cases, or with a high statistical probability. This statement has the general character of a law of statistical form, and while the probability value is not specified, the statement indicates that it is high. But in contrast to the cases of deductive-nomological and deductive-statistical explanation, the explanans consisting of this statistical law together with the statement that the patient did receive penicillin obviously does not imply the explanandum statement, 'the patient recovered', with deductive certainty, but only, as we might say, with high likelihood, or near certainty. Briefly, then, the explanation amounts to this argument:

1*a* The particular case of illness of John Jones—let us call it *j*—was an instance of severe streptococcal infection (Sj) which was treated with large doses of penicillin (Pj); and the statistical probability $p(R, S \cdot P)$ of recovery in cases where S and P are present close to 1; hence, the case was practically certain to end in recovery (Rj).*

FROM *Aspects of Scientific Explanation* (New York: Free Press, 1965), 381–83, 394–403.

* Throughout this paper, Hempel uses a dot to stand for conjunction, a bar over a letter to stand for negation, and a comma within parentheses to represent conditional probabilities. Thus, for example, $p(R, S \cdot \bar{P})$ means the probability of R given S and not-P.

This argument might invite the following schematization:

1*b* $p(R, S \cdot P)$ is close to 1
$Sj \cdot Pj$

(Therefore:) It is practically certain (very likely) that Rj

In the literature on inductive inference, arguments thus based on statistical hypotheses have often been construed as having this form or a similar one. On this construal, the conclusion characteristically contains a modal qualifier such as 'almost certainly,' 'with high probability', 'very likely', etc. But the conception of arguments having this character is untenable. For phrases of the form 'it is practically certain that *p*' or 'It is very likely that *p*', where the place of '*p*' is taken by some statement, are not complete self-contained sentences that can be qualified as either true or false. The statement that takes the place of '*p*'—for example, 'Rj'—is either true or false, quite independently of whatever relevant evidence may be available, but it can be qualified as more or less likely, probable, certain, or the like only *relative to some body of evidence*. One and the same statement, such as 'Rj', will be certain, very likely, not very likely, highly likely, and so forth, depending upon what evidence is considered. The phrase 'it is almost certain that Rj' taken by itself is therefore neither true nor false; and it cannot be inferred from the premises specified in (1*b*) nor from any other statements.

The confusion underlying the schematization (1*b*) might be further illuminated by considering its analogue for the case of deductive arguments. The force of a deductive inference, such as that from 'all F are G' and 'a is F' to 'a is G', is sometimes indicated by saying that if the premises are true, then the conclusion is necessarily true or is certain to be true—a phrasing that might suggest the schematization

All F are G
a is F

(Therefore:) It is necessary (certain) that a is G

But clearly the given premises—which might be, for example, 'all men are mortal' and 'Socrates is a man'—do not establish the sentence 'a is G' ('Socrates is mortal') as a necessary or certain truth. The certainty referred to in the informal paraphrase of the argument is relational: the statement 'a is G' is certain, or necessary, *relative to the specified premises*; i.e., their truth will guarantee its truth—which means nothing more than that 'a is G' is a logical consequence of those premises.

Analogously, to present our statistical explanation in the manner of schema (1*b*) is to misconstrue the function of the words 'almost certain' or 'very likely' as they occur in the formal wording of the explanation.

Those words clearly must be taken to indicate that on the evidence provided by the explanans, or relative to that evidence, the explanandum is practically certain or very likely, i.e., that

1c 'Rj' is practically certain (very likely) relative to the explanans containing the sentences '$p(R, S \cdot P)$ is close to 1' and '$Sj \cdot Pj$'.[1]

The explanatory argument misrepresented by (1*b*) might therefore suitably be schematized as follows:

1*d* $p(R, S \cdot P)$ is close to 1
$Sj \cdot Pj$
——— [makes practically certain]
Rj

In this schema, the double line separating the "premises" from the "conclusion" is to signify that the relation of the former to the latter is not that of deductive implication but that of inductive support, the strength of which is indicated in square brackets.[2] . . .

2 | The Problem of Explanatory Ambiguity

Consider once more the explanation (1*d*) of recovery in the particular case j of John Jones's illness. The statistical law there invoked claims recovery in response to penicillin only for a high percentage of streptococcal infections, but not for all of them; and in fact, certain streptococcus strains are resistant to penicillin. Let us say that an occurrence, e.g. a particular case of illness, has the property S^* (or belongs to the class S^*) if it is an instance of infection with a penicillin-resistant streptococcus strain. Then the probability of recovery among randomly chosen instances of S^* which are treated with penicillin will be quite small, i.e, $p(R, S^* \cdot P)$ will be close to 0 and the probability of nonrecovery, $p(\bar{R}, S^* \cdot P)$ will be close to 1. But suppose now that Jones's illness is in fact a streptococcal infection of the penicillin-resistant variety, and consider the following argument:

2*a* $p(\bar{R}, S^* \cdot P)$ is close to 1
$S^*j \cdot Pj$
——— [makes practically certain]
$\bar{R}j$

This "rival" argument has the same form as (1*d*), and on our assumptions its premises are true, just like those of (1*d*). Yet its conclusion is the contradictory of the conclusion of (1*d*).

Or suppose that Jones is an octogenarian with a weak heart, and that

in this group, S^{**}, the probability of recovery from a streptococcus infection in response to penicillin treatment, $p(R, S^{**} \cdot P)$, is quite small. Then, there is the following rival argument to (1*d*), which presents Jones's nonrecovery as practically certain in the light of premises which are true:

2*b* $p(\bar{R}, S^{**} \cdot P)$ is close to 1
$S^{**}j \cdot Pj$
——— [makes practically certain]
$\bar{R}j$

The peculiar logical phenomenon here illustrated will be called the *ambiguity of inductive-statistical explanation* or, briefly, of *statistical explanation*. This ambiguity derives from the fact that a given individual event (e.g., Jones's illness) will often be obtainable by random selection from any one of several "reference classes" (such as $S \cdot P$, $S^* \cdot P$, $S^{**} \cdot P$), with respect to which the kind of occurrence (e.g., R) instantiated by the given event has very different statistical probabilities. Hence, for a proposed probabilistic explanation with true explanans which confers near certainty upon a particular event, there will often exist a rival argument of the same probabilistic form and with equally true premises which confers near certainty upon the nonoccurrence of the same event. And any statistical explanation for the occurrence of an event must seem suspect if there is the possibility of a logically and empirically equally sound probabilistic account for its nonoccurrence. *This predicament has no analogue in the case of deductive explanation*; for if the premises of a proposed deductive explanation are true then so is its conclusion; and its contradictory, being false, cannot be a logical consequence of a rival set of premises that are equally true.

Here is another example of the ambiguity of I-S explanation: Upon expressing surprise at finding the weather in Stanford warm and sunny on a date as autumnal as November 27, I might be told, by way of explanation, that this was rather to be expected because the probability of warm and sunny weather (W) on a November day in Stanford (N) is, say, .95. Schematically, this account would take the following form, where 'n' stands for 'November 27':

2*c* $p(W, N) = .95$
Nn
——— [.95]
Wn

But suppose it happens to be the case that the day before, November 26, was cold and rainy, and that the probability for the immediate successors (S) of cold and rainy days in Stanford to be warm and sunny is .2; then the account (2*c*) has a rival in the following argument which,

by reference to equally true premises, presents it as fairly certain that November 27 is not warm and sunny:

$$2d \quad \frac{p(\overline{W}, S) = .8 \qquad Sn}{\overline{W}n} \; [.8]$$

In this form, the problem of ambiguity concerns I-S arguments whose premises are in fact true, no matter whether we are aware of this or not. But, as will now be shown, the problem has a variant that concerns explanations whose explanans statements, no matter whether in fact true or not, are *asserted or accepted* by empirical science at the time when the explanation is proffered or contemplated. This variant will be called *the problem of the epistemic ambiguity of statistical explanation*, since it refers to what is presumed to be known in science rather than to what, perhaps unknown to anyone, is in fact the case.

Let K_t be the class of all statements asserted or accepted by empirical science at time t. This class then represents the total scientific information, or "scientific knowledge" at time t. The word 'knowledge' is here used in the sense in which we commonly speak of the scientific knowledge at a given time. It is not meant to convey the claim that the elements of K_t are true, and hence neither that they are definitely known to be true. No such claim can justifiably be made for any of the statements established by empirical science; and the basic standards of scientific inquiry demand that an empirical statement, however well supported, be accepted and thus admitted to membership in K_t only tentatively, i.e., with the understanding that the privilege may be withdrawn if unfavorable evidence should be discovered. The membership of K_t therefore changes in the course of time; for as a result of continuing research, new statements are admitted into that class; others may come to be discredited and dropped. Henceforth, the class of accepted statements will be referred to simply as K when specific reference to the time in question is not required. We will assume that K is logically consistent and that it is closed under logical implication, i.e., that it contains every statement that is logically implied by any of its subsets.

The *epistemic ambiguity of I-S explanation* can now be characterized as follows: The total set K of accepted scientific statements contains different subsets of statements which can be used as premises in arguments of the probabilistic form just considered, and which confer high probabilities on logically contradictory "conclusions." Our earlier examples (2*a*), (2*b*) and (2*c*), (2*d*) illustrate this point if we assume that the premises of those arguments all belong to K rather than that they are all true. If one of two such rival arguments with premises in K is proposed as an explanation of an event considered, or acknowledged, in science to have occurred, then the conclusion of the argument, i.e., the explanandum statement, will accordingly belong to K as well. And since K is consistent, the conclusion of the rival argument will not belong to K. Nonetheless it is disquieting that we should be able to say: No matter whether we are informed that the event in question (e.g. warm and sunny weather on November 27 in Stanford) did occur or that it did not occur, we can produce an explanation of the reported outcome in either case; and an explanation, moreover, whose premises are scientifically established statements that confer a high logical probability upon the reported outcome.

This epistemic ambiguity, again, has no analogue for deductive explanation; for since K is logically consistent, it cannot contain premise-sets that imply logically contradictory conclusions.

Epistemic ambiguity also bedevils the predictive use of statistical arguments. Here, it has the alarming aspect of presenting us with two rival arguments whose premises are scientifically well established, but one of which characterizes a contemplated future occurrence as practically certain, whereas the other characterizes it as practically impossible. Which of such conflicting arguments, if any, are rationally to be relied on for explanation or for prediction?

3 | The Requirement of Maximal Specificity and the Epistemic Relativity of Inductive-Statistical Explanation

Our illustrations of explanatory ambiguity suggest that a decision on the acceptability of a proposed probabilistic explanation or prediction will have to be made in the light of all the relevant information at our disposal. This is indicated also by a general principle whose importance for inductive reasoning has been acknowledged, if not always very explicitly, by many writers, and which has recently been strongly emphasized by Carnap, who calls it *the requirement of total evidence*. Carnap formulates it as follows: "in the application of inductive logic to a given knowledge situation, the total evidence available must be taken as basis for determining the degree of confirmation."[3] Using only a part of the total evidence is permissible if the balance of the evidence is irrelevant to the inductive "conclusion," i.e., if on the partial evidence alone, the conclusion has the same confirmation, or logical probability, as on the total evidence.[4]

The requirement of total evidence is not a postulate nor a theorem of inductive logic; it is not concerned with the formal validity of inductive arguments. Rather, as Carnap has stressed, it is a maxim for the *application* of inductive logic; we might say that it states a necessary condition of

rationality of any such application in a given "knowledge situation," which we will think of as represented by the set K of all statements accepted in the situation.

But in what manner should the basic idea of this requirement be brought to bear upon probabilistic explanation? Surely we should not insist that the explanans must contain all and only the empirical information available at the time. Not *all* the available information, because otherwise all probabilistic explanations acceptable at time t would have to have the same explanans, K_t; and not *only* the available information, because a proffered explanation may meet the intent of the requirement in not overlooking any relevant information available, and may nevertheless invoke some explanans statements which have not as yet been sufficiently tested to be included in K_t.

The extent to which the requirement of total evidence should be imposed upon statistical explanations is suggested by considerations such as the following. A proffered explanation of Jones's recovery based on the information that Jones had a streptococcal infection and was treated with penicillin, and that the statistical probability for recovery in such cases is very high, is unacceptable if K includes the further information that Jones's streptococci were resistant to penicillin, or that Jones was an octogenarian with a weak heart, and that in these reference classes the probability of recovery is small. Indeed, one would want an acceptable explanation to be based on a statistical probability statement pertaining to the narrowest reference class of which, according to our total information, the particular occurrence under consideration is a member. Thus, if K tells us not only that Jones had a streptococcus infection and was treated with penicillin, but also that he was an octogenarian with a weak heart (and if K provides no information more specific than that) then we would require that an acceptable explanation of Jones's response to the treatment be based on a statistical law stating the probability of that response in the narrowest reference class to which our total information assigns Jones's illness, i.e., the class of streptococcal infections suffered by octogenarians with weak hearts.[5]

Let me amplify this suggestion by reference to an example concerning the use of the law that the half-life of radon is 3.82 days in accounting for the fact that the residual amount of radon to which a sample of 10 milligrams was reduced in 7.64 days was within the range from 2.4 to 2.6 milligrams. According to present scientific knowledge, the rate of decay of a radioactive element depends solely upon its atomic structure as characterized by its atomic number and its mass number, and it is thus unaffected by the age of the sample and by such factors as temperature, pressure, magnetic and electric forces, and chemical interactions. Thus, by specifying the half-life of radon as well as the initial mass of the sample and the time interval in question, the explanans takes into account all the

available information that is relevant to appraising the probability of the given outcome by means of statistical laws. To state the point somewhat differently: Under the circumstances here assumed, our total information K assigns the case under study first of all to the reference class say F_1, of cases where a 10 milligram sample of radon is allowed to decay for 7.64 days; and the half-life law for radon assigns a very high probability, within F_1, to the "outcome," say G, consisting in the fact that the residual mass of radon lies between 2.4 and 2.6 milligrams. Suppose now that K also contains information about the temperature of the given sample, the pressure and relative humidity under which it is kept, the surrounding electric and magnetic conditions, and so forth, so that K assigns the given case to a reference class much narrower than F_1, let us say, $F_1F_2F_3 \ldots F_n$. Now the theory of radioactive decay, which is equally included in K, tells us that the statistical probability of G within this narrower class is the same as within G. For this reason, it suffices in our explanation to rely on the probability $p(G, F_1)$.

Let us note, however, that "knowledge situations" are conceivable in which the same argument would not be an acceptable explanation. Suppose, for example, that in the case of the radon sample under study, the amount remaining one hour before the end of the 7.64 day period happens to have been measured and found to be 2.7 milligrams, and thus markedly in excess of 2.6 milligrams—an occurrence which, considering the decay law for radon, is highly improbable, but not impossible. That finding, which then forms part of the total evidence K, assigns the particular case at hand to a reference class, say F^*, within which, according to the decay law for radon, the outcome G is highly improbable since it would require a quite unusual spurt in the decay of the given sample to reduce the 2.7 milligrams, within the one final hour of the test, to an amount falling between 2.4 and 2.6 milligrams. Hence, the additional information here considered may not be disregarded, and an explanation of the observed outcome will be acceptable only if it takes account of the probability of G in the narrower reference class, i.e., $p(G, F_1F^*)$. (The theory of radioactive decay implies that this probability equals $p(G, F^*)$, so that as a consequence the membership of the given case in F_1 need not be explicitly taken into account.)

The requirement suggested by the preceding considerations can now be stated more explicitly; we will call it the *requirement of maximal specificity for inductive-statistical explanations*. Consider a proposed explanation of the basic statistical form

$$3a \quad \begin{array}{l} p(G, F) = r \\ \hline Fb \\ \hline\hline Gb \end{array} \; [r]$$

Let s be the conjunction of the premises, and, if K is the set of all statements accepted at the given time, let k be a sentence that is logically equivalent to K (in the sense that k is implied by K and in turn implies every sentence in K). Then, to be rationally acceptable in the knowledge situation represented by K, the proposed explanation (3a) must meet the following condition (the requirement of maximal specificity): If $s \cdot k$ implies[6] that b belongs to a class F_1, and that F_1 is a subclass of F, then $s \cdot k$ must also imply a statement specifying the statistical probability of G in F_1, say

$$p(G, F_1) = r_1$$

Here, r_1 must equal r unless the probability statement just cited is simply a theorem of mathematical probability theory.

The qualifying unless-clause here appended is quite proper, and its omission would result in undesirable consequences. It is proper because theorems of pure mathematical probability theory cannot provide an explanation of empirical subject matter. They may therefore be discounted when we inquire whether $s \cdot k$ might not give us statistical laws specifying the probability of G in reference classes narrower than F. And the omission of the clause would prove troublesome, for if (3a) is proffered as an explanation, then it is presumably accepted as a fact that Gb; hence 'Gb' belongs to K. Thus K assigns b to the narrower class $F \cdot G$, and concerning the probability of G in that class, $s \cdot k$ trivially implies the statement that $p(G, F \cdot G) = 1$, which is simply a consequence of the measure-theoretical postulates for statistical probability. Since $s \cdot k$ thus implies a more specific probability statement for G than that invoked in (3a), the requirement of maximal specificity would be violated by (3a)— and analogously by any proffered statistical explanation of an event that we take to have occurred—were it not for the unless-clause, which, in effect, disqualifies the notion that the statement '$p(G, F \cdot G) = 1$' affords a more appropriate law to account for the presumed fact that Gb.

The requirement of maximal specificity, then, is here tentatively put forward as characterizing the extent to which the requirement of total evidence properly applies to inductive-statistical explanations. The general idea thus suggested comes to this: In formulating or appraising an I-S explanation, we should take into account all that information provided by K which is of potential *explanatory* relevance to the explanandum event, i.e., all pertinent statistical laws, and such particular facts as might be connected, by the statistical laws, with the explanandum event.[7]

The requirement of maximal specificity disposes of the problem of epistemic ambiguity; for it is readily seen that of two rival statistical arguments with high associated probabilities and with premises that all belong to K, at least one violates the requirement of maximum specificity. Indeed, let

$$\frac{p(G, F) = r_1 \quad Fb}{Gb}\ [r_1] \qquad \text{and} \qquad \frac{p(\bar{G}, H) = r_2 \quad Hb}{\bar{G}b}\ [r_2]$$

be the arguments in question, with r_1 and r_2 close to 1. Then, since K contains the premises of both arguments, it assigns b to both F and H and hence to $F \cdot H$. Hence if both arguments satisfy the requirement of maximal specificity, K must imply that

$$\begin{aligned} & p(G, F \cdot H) = p(G, F) = r_1 \\ & p(\bar{G}, F \cdot H) = p(\bar{G}, H) = r_2 \\ \text{But} \quad & p(G, F \cdot H) + p(\bar{G}, F \cdot H) = 1 \\ \text{Hence} \quad & r_1 + r_2 = 1 \end{aligned}$$

and this is an arithmetic falsehood, since r_1 and r_2 are both close to 1; hence it cannot be implied by the consistent class K.

Thus, for I-S explanations that meet the requirement of maximal specificity the problem of epistemic ambiguity no longer arises. We are *never* in a position to say: No matter whether this particular event did or did not occur, we can produce an acceptable explanation of either outcome; and an explanation, moreover, whose premises are scientifically accepted statements which confer a high logical probability upon the given outcome.

While the problem of epistemic ambiguity has thus been resolved, ambiguity in the first sense discussed [in section 2] remains unaffected by our requirement; i.e., it remains the case that for a given statistical argument with true premises and a high associated probability, there may exist a rival one with equally true premises and with a high associated probability, whose conclusion contradicts that of the first argument. And though the set K of statements accepted at any time never includes all statements that are in fact true (and no doubt many that are false), it is perfectly possible that K should contain the premises of two such conflicting arguments; but as we have seen, at least one of the latter will fail to be rationally acceptable because it violates the requirement of maximal specificity.

The preceding considerations show that *the concept of statistical explanation for particular events is essentially relative to a given knowledge situation as represented by a class K of accepted statements.* Indeed, the requirement of maximal specificity makes explicit and unavoidable reference to such a class, and it thus serves to characterize the concept of "I-S explanation relative to the knowledge situation represented by K." We will refer to this characteristic as the *epistemic relativity of statistical explanation.*

It might seem that the concept of deductive explanation possesses the

same kind of relativity, since whether a proposed D-N or D-S [deductive-statistical] account is acceptable will depend not only on whether it is deductively valid and makes essential use of the proper type of general law, but also on whether its premises are well supported by the relevant evidence at hand. Quite so; and this condition of empirical confirmation applies equally to statistical explanations that are to be acceptable in a given knowledge situation. But the epistemic relativity that the requirement of maximal specificity implies for I-S explanations is of quite a different kind and has no analogue for D-N explanations. For the specificity requirement is not concerned with the evidential support that the total evidence K affords for the explanans statements: it does not demand that the latter be included in K, nor even that K supply supporting evidence for them. It rather concerns what may be called the concept of a *potential* statistical explanation. For it stipulates that no matter how much evidential support there may be for the explanans, a proposed I-S explanation is not acceptable if its potential explanatory force with respect to the specified explanandum is vitiated by statistical laws which are included in K but not in the explanans, and which might permit the production of rival statistical arguments. As we have seen, this danger never arises for deductive explanations. Hence, these are not subject to any such restrictive condition, and the notion of a potential deductive explanation (as contradistinguished from a deductive explanation with well-confirmed explanans) requires no relativization with respect to K.

As a consequence, we can significantly speak of true D-N and D-S explanations: they are those potential D-N and D-S explanations whose premises (and hence also conclusions) are true—no matter whether this happens to be known or believed, and thus no matter whether the premises are included in K. But this idea has no significant analogue for I-S explanation since, as we have seen, the concept of potential statistical explanation requires relativization with respect to K.

■ | Notes

1. Phrases such as 'It is almost certain (very likely) that j recovers', even when given the relational construal here suggested, are ostensibly concerned with relations between propositions, such as those expressed by the sentences forming the conclusion and the premises of an argument. For the purpose of the present discussion, however, involvement with propositions can be avoided by construing the phrases in question as expressing logical relations between corresponding *sentences*, e.g., the conclusion-sentence and the premise-sentence of an argument. This construal, which underlies the formulation of (1c), will be adopted in this essay, though for the sake of convenience we may occasionally use a paraphrase.

2. In the familiar schematization of deductive arguments, with a single line separating the premises from the conclusion, no explicit distinction is made between a weaker and a stronger claim, either of which might be intended; namely (i) that the premises logically imply the conclusion and (ii) that, in addition, the premises are true. In the case of our probabilistic argument, (1c) expresses a weaker claim, analogous to (i), whereas (1d) may be taken to express a "proffered explanation" (the term is borrowed from I. Scheffler, 'Explanation, Prediction, and Abstraction', *British Journal for the Philosophy of Science* 7 (1957), sect. 1) in which, in addition, the explanatory premises are—however tentatively—asserted as true.

The considerations here outlined concerning the use of terms like 'probably' and 'certainly' as modal qualifiers of individual statements seem to me to militate also against the notion of categorical probability statement that C. I. Lewis sets forth in the following passage (italics the author's):

> Just as 'If D then (certainly) P, and D is the fact', leads to the categorical consequence, 'Therefore (certainly) P'; so too, 'If D then probably P, and D is the fact', leads to a categorical consequence expressed by 'It is probable that P'. And this conclusion is not merely the statement over again of the probability relation between 'P' and 'D'; any more than 'Therefore (certainly) P' is the statement over again of 'If D then (certainly) P'. 'If the barometer is high, tomorrow will probably be fair; and the barometer *is* high', categorically assures something expressed by 'Tomorrow will probably be fair'. This probability is still relative to the grounds of judgment; but if these grounds are actual, and contain all the available evidence which is pertinent, then it is not only categorical but may fairly be called *the* probability of the event in question (1946: 319).

This position seems to me to be open to just those objections suggested in the main text. If 'P' is a statement, then the expressions 'certainly P' and 'probably P' as envisaged in the quoted passage are not statements. If we ask how one would go about trying to ascertain whether they were true, we realize that we are entirely at a loss unless and until a reference set of statements or assumptions has been specified relative to which P may then be found to be certain, or to be highly probable, or neither. The expressions in question, then, are essentially incomplete; they are elliptic formulations of relational statements; neither of them can be the conclusion of an inference. However plausible Lewis's suggestion may seem, there is no analogue in inductive logic to *modus ponens*, or the "rule of detachment," of deductive logic, which, given the information that 'D' and also 'if D then P', are true statements, authorizes us to detach the consequent 'P' in the conditional premise and to assert it as a self-contained statement which must then be true as well.

At the end of the quoted passage, Lewis suggests the important idea that 'probably P' might be taken to mean that the total relevant evidence available at the time confers high probability upon P. But even this statement is relational in that it tacitly refers to some unspecified time, and, besides, his general notion of a categorical probability statement as a conclusion of an argument is not made dependent on the assumption that the premises of the argument include all the relevant evidence available.

It must be stressed, however, that elsewhere in his discussion, Lewis emphasizes the relativity of (logical) probability, and, thus, the very characteristic that rules out the conception of categorical probability statements.

Similar objections apply, I think, to Toulmin's construal of probabilistic arguments; cf. Toulmin (1958) and the discussion in Hempel (1960), sects. 1–3.

3. R. Carnap, *Logical Foundations of Probability* (Chicago, 1950), 211. The requirement is suggested, e.g., in the passage from Lewis quoted in n. [2]. Similarly Williams speaks of "the most fundamental of all rules of probability logic, that 'the' probability of any proposition is its probability in relation to the known premises and them only" (*The Ground of Induction* (Cambridge, Mass., 1947), 72).

I am greatly indebted to Professor Carnap for having pointed out to me in 1945, when I first noticed the ambiguity of probabilistic arguments, that this was but one of several apparent paradoxes of inductive logic that result from disregard of the requirement of total evidence.

S. F. Barker, *Induction and Hypothesis* (Ithaca, NY, 1957), 70–78, has given a lucid independent presentation of the basic ambiguity of probabilistic arguments, and a skeptical appraisal of the requirement of total evidence as a means of dealing with the problem. However, I will presently suggest a way of remedying the ambiguity of probabilistic explanation with the help of a rather severely modified version of the requirement of total evidence. It will be called the requirement of maximal specificity, and is not open to the same criticism.

4. Cf. Carnap, *Logical Foundations*, 211 and 494.

5. This idea is closely related to one used by H. Reichenbach, (cf. *The Theory of Probability* (Berkeley, Calif., and Los Angeles, 1949), sect. 72) in an attempt to show that it is possible to assign probabilities to individual events within the framework of a strictly statistical conception of probability. Reichenbach proposed that the probability of a single event, such as the safe completion of a particular scheduled flight of a given commercial plane, be construed as the statistical probability which the *kind* of event considered (safe completion of a flight) possesses within the narrowest reference class to which the given case (the specified flight of the given plane) belongs, and for which reliable statistical information is available (e.g., the class of scheduled flights undertaken so far by planes of the line to which the given plane belongs, and under weather conditions similar to those prevailing at the time of the flight in question).

6. Reference to $s \cdot k$ rather than to k is called for because, as was noted earlier, we do not construe the condition here under discussion as requiring that all the explanans statements invoked be scientifically accepted at the time in question, and thus be included in the corresponding class K.

7. By its reliance on this general idea, and specifically on the requirement of maximal specificity, the method here suggested for eliminating the epistemic ambiguity of statistical explanation differs substantially from the way in which I attempted in an earlier study (Hempel, 'Deductive-Nomological vs. Statistical Explanation', esp. sect. 10) to deal with the same problem. In that study, which did not distinguish explicitly between the two types of explanatory ambiguity characterized earlier in this section, I applied the requirement of total evidence to statistical explanations in a manner which presupposed that the explanans of any acceptable explanation belongs to the class K, and which then demanded that the probability which the explanans confers upon the explanandum be equal to that which the total evidence, K, imparts to the explanandum. The reasons why this approach seems unsatisfactory to me are suggested by the arguments set forth in the present section. Note in particular that, if strictly enforced, the requirement of total evidence would preclude the possibility of any significant statistical explanation for events whose occurrence is regarded as an established fact in science: for any sentence describing such an occurrence is logically implied by K and thus trivially has the logical probability 1 relative to K.

References

Barker, S. F., *Induction and Hypothesis* (Ithaca, NY: Cornell University Press, 1957).

Carnap, R., *Logical Foundations of Probability* (Chicago: University of Chicago Press 1950; second, rev., edn. 1962).

Feigl, H., and Maxwell, G. (eds.), *Minnesota Studies in the Philosophy of Science*, iii (Minneapolis: University of Minnesota Press, 1962).

Hempel, C. G., 'Inductive Inconsistencies', in *Synthese*, 12 (1960), 439–69, repr. in *Aspects of Scientific Explanation* (New York and London: Free Press and Collier Macmillan, 1965), 53–79.

——'Deductive-Nomological vs. Statistical Explanation', in Feigl and Maxwell (eds.), 98–169.

Lewis, C. I., *An Analysis of Knowledge and Valuation* (La Salle, Ill.: Open Court, 1946).

Reichenbach, H., *The Theory of Probability* (Berkeley, Calif., and Los Angeles: The University of California Press, 1949).

Scheffler, I., 'Explanation, Prediction, and Abstraction', *The British Journal for the Philosophy of Science* 7 (1957) 293–309.

Toulmin, S., *The Uses of Argument* (Cambridge: Cambridge University Press, 1958).

Williams, D. C., *The Ground of Induction* (Cambridge, Mass.: Harvard University Press, 1947).