CARL G. HEMPEL

Deductive-Nomological vs. Statistical Explanation

1. *Objectives of This Essay.*

This essay is concerned with the form and function of explanation in the sense in which it is sought, and often achieved, by empirical science. It does not propose to examine all aspects of scientific explanation; in particular, a closer study of historical explanation falls outside the purview of the present investigation. My main object is to propose, and to elaborate to some extent, a distinction of two basic modes of explanation—and similarly of prediction and retrodiction—which will be called the deductive and the inductive mode.[1]

The structure of deductive explanation and prediction conforms to what is now often called the covering-law model: it consists in the deduction of whatever is being explained or predicted from general laws in conjunction with information about particular facts. The logic of this procedure was examined in some earlier articles of mine, and especially in a study carried out in collaboration with P. Oppenheim.[2]

Since then, various critical comments and constructive suggestions concerning those earlier efforts have appeared in print, and these as well

[1] This distinction was developed briefly in Hempel [25], Sec. 2.

[2] See Hempel [24], especially Secs. 1–4; Hempel [25]; and Hempel and Oppenheim [26]. This latter article will henceforth be referred to as *SLE*. The point of these discussions was to give a more precise and explicit statement of the deductive model of scientific explanation and to exhibit and analyze some of the logical and methodological problems to which it gives rise: the general conception of explanation as deductive subsumption under more general principles had been set forth much earlier by a variety of authors, some of whom are listed in *SLE*, fn. 4. In fact, in 1934 that conception was explicitly presented in the following passage of an introductory textbook: "Scientific explanation consists in subsuming under some rule or law which expresses an invariant character of a group of events, the particular event it is said to explain. Laws themselves may be explained, and in the same manner, by showing that they are consequences of more comprehensive theories." (Cohen and Nagel [10], p. 397.) The conception of the explanation of laws by deduction from theories was

as discussions with interested friends and with my students have led me to reconsider the basic issues concerning the deductive model of scientific explanation and prediction. In the first of the two principal parts of this essay, I propose to give a brief survey of those issues, to modify in certain respects the ideas set forth in the earlier articles, and to examine some new questions concerning deductive explanation, deductive prediction, and related procedures.

The second major part of the present study is an attempt to point out, and to shed some light on, certain fundamental problems in the logic of inductive explanation and prediction.

Part I. Deductive-Nomological Systematization

2. *The Covering-Law Model of Explanation.*

The deductive conception of explanation is suggested by cases such as the following: The metal screwtop on a glass jar is tightly stuck; after being placed in warm water for a short while, it can be readily removed. The familiar explanation of this phenomenon is, briefly, to the effect that the metal has a higher coefficient of thermal expansion than glass, so that a given rise in temperature will produce a larger expansion of the lid than of the neck of the glass jar; and that, in addition, though the metal is a good conductor of heat, the temperature of the lid will temporarily be higher than that of the glass—a fact which further increases the difference between the two perimeters. Thus, the loosening of the lid is here explained by showing that it came about, by virtue of certain antecedent circumstances, in accordance with certain physical laws. The explanation may be construed as an argument in which the occurrence of the event in question is inferred from information expressed by statements of two kinds: (a) general laws, such as those concerning the thermal conductivity of metal and the coefficients of expansion for metal and for glass, as well as the law that heat will be transferred from one body to another of lower temperature with which it is in contact; (b) statements describing particular circumstances, such as that the jar is made of glass, the lid of metal; that initially, at room

developed in great detail by N. R. Campbell; for an elementary account see his book [4], which was first published in 1921. K. R. Popper, too, has set forth this deductive conception of explanation in several of his publications (cf. fn. 4 in *SLE*); his earliest statement appears in Sec. 12 of his book [38], which has at long last been published in a considerably expanded English version [40].

temperature, the lid fitted very tightly on the top of the jar; and that then the top with the lid on it was immersed in hot water. To show that the loosening of the lid occurred "by virtue of" the circumstances in question, and "in accordance with" those laws, is then to show that the statement describing the result can be validly inferred from the specified set of premises.

Thus construed, the explanation at hand is a deductive argument of this form:

$$(2.1) \quad \frac{\begin{array}{c} L_1, L_2, \ldots, L_r \\ C_1, C_2, \ldots, C_k \end{array}}{E}$$

Here, $L_1, L_2, \ldots, L_r$ are general laws and $C_1, C_2, \ldots, C_k$ are statements of particular occurrences, facts, or events; jointly, these premises form the explanans. The conclusion E is the explanandum statement; it describes the phenomenon (or event, etc.) to be explained, which will also be called the explanandum phenomenon (or event, etc.); thus, the word 'explanandum' will be used to refer ambiguously either to the explanandum statement or to the explanandum phenomenon. Inasmuch as the sentence E is assumed to be a logical consequence of the premises, an explanatory argument of form (2.1) deductively subsumes the explanandum under *"covering laws."*[3] I will say, therefore, that (2.1) represents the *covering-law model of explanation*. More specifically, I will refer to explanatory arguments of the form (2.1) as *deductive-nomological*, or briefly as deductive, explanations: as will be shown later, there are other explanations invoking general laws that will have to be construed as inductive rather than as deductive arguments.

In my illustration, the explanandum is a particular event, the loosening of a certain lid, which occurs at a definite place and time. But deductive subsumption under general laws can serve also to explain general uniformities, such as those asserted by laws of nature. For example, the

[3] The suggestive terms 'covering law' and 'covering-law model' are borrowed from Dray, who, in his book [13], presents a lucid and stimulating critical discussion of the question whether, or to what extent, historical explanation conforms to the deductive pattern here considered. To counter a misunderstanding that might be suggested by some passages in Ch. II, Sec. 1 of Dray's book, I would like to emphasize that the covering-law model must be understood as permitting reference to *any number of laws* in the explanation of a given phenomenon: there should be no restriction to just *one* "covering law" in each case.

uniformity expressed by Galileo's law for free fall can be explained by deduction from the general laws of mechanics and Newton's law of gravitation, in conjunction with statements specifying the mass and radius of the earth. Similarly, the uniformities expressed by the law of geometrical optics can be explained by deductive subsumption under the principles of the wave theory of light.[4]

3. *Truth and Confirmation of Deductive Explanations.*

In *SLE* (Section 3) two basic requirements are imposed upon a scientific explanation of the deductive-nomological variety:[5] (i) It must be a deductively valid argument of the form (2.1), whose premises include at least one general law essentially, i.e., in such a way that if the law were deleted, the argument would no longer be valid. Intuitively, this means that reliance on general laws is essential to this type of explanation; a given phenomenon is here explained, or accounted for, by showing that it conforms to a general nomic pattern. (ii) The sentences constituting the explanans must be true, and hence so must the explanandum sentence. This second requirement was defended by the following consideration: suppose we required instead that the explanans be highly

[4] More accurately, the explanation of a general law by means of a theory will usually show (1) that the law holds only within a certain range of application, which may not have been made explicit in its standard formulation; (2) that even within that range, the law holds only in close approximation, but not strictly. This point is well illustrated by Duhem's emphatic reminder that Newton's law of gravitation, far from being an inductive generalization of Kepler's laws, is actually incompatible with them, and that the credentials of Newton's theory lie rather in its enabling us to compute the perturbations of the planets, and thus their deviations from the orbits assigned to them by Kepler. (See Duhem [14], pp. 312ff, and especially p. 317. The passages referred to here are included in the excerpts from P. P. Wiener's translation of Duhem's work that are reprinted in Feigl and Brodbeck [15], under the title "Physical Theory and Experiment.")

Analogously, Newtonian theory implies that the acceleration of a body falling freely in a vacuum toward the earth will increase steadily, though over short distances it will be very nearly constant. Thus, strictly speaking, the theory contradicts Galileo's law, but shows the latter to hold true in very close approximation within a certain range of application. A similar relation obtains between the principles of wave optics and those of geometrical optics.

[5] No claim was made that this is the only kind of scientific explanation; on the contrary, at the end of Sec. 3, it was emphasized that "Certain cases of scientific explanation involve 'subsumption' of the explanandum under a set of laws of which at least some are statistical in character. Analysis of the peculiar logical structure of that type of subsumption involves difficult special problems. The present essay will be restricted to an examination of the causal type of explanation . . ." A similar explicit statement is included in the final paragraph of Sec. 7 and in Sec. 5.3 of the earlier article, Hempel [24]. These passages seem to have been overlooked by some critics of the covering-law model.

confirmed by all the relevant evidence available, though it need not necessarily be true. Now it might happen that the explanans of a given argument of the form (2.1) was well confirmed at a certain earlier stage of scientific research, but strongly disconfirmed by the more comprehensive evidence available at a later time, say, the present. In this event, we would have to say that the explanandum was correctly explained by the given argument at the earlier stage, but not at the later one. And this seemed counterintuitive, for common usage appeared to construe the correctness of a given explanation as no more time dependent than, say, the truth of a given statement. But this justification, with its reliance on a notion of correctness that does not appear in the proposed definition of explanation, is surely of questionable merit. For in reference to explanations as well as in reference to statements, the vague idea of correctness can be construed in two different ways, both of which are of interest and importance for the logical analysis of science: namely, as truth in the semantical sense, which is independent of any reference to time or to evidence; or as confirmation by the available relevant evidence—a concept which is clearly time dependent. We will therefore distinguish between *true explanations*, which meet the requirement of truth for their explanans, and *explanations that are more or less well confirmed* by a given body of evidence (e.g., by the total evidence available). These two concepts can be introduced as follows:

First, we define a *potential explanation* (of deductive-nomological form)[6] as an argument of the form (2.1) which meets all the requirements indicated earlier, except that the statements forming its explanans and explanandum need not be true. But the explanans must still contain a set of sentences, $L_1, L_2, \ldots, L_r$, which are lawlike, i.e., which are like laws except for possibly being false.[7] Sentences of this kind will also be called *nomic*, or *nomological, statements*. It is this notion of potential

[6] This was done already in *SLE*, Sec. 7.

[7] The term 'lawlike sentence' and the general characterization given here of its intended meaning are from Goodman [20]. The difficult problem of giving an adequate general characterization of those sentences which if true would constitute laws will not be dealt with in the present essay. For a discussion of the issues involved, see, for example, *SLE*, Secs. 6–7; Braithwaite [3], Ch. IX, where the central question is described as concerning "the nature of the difference, if any, between 'nomic laws' and 'mere generalizations'"; and the new inquiry into the subject by Goodman [20, 21]. All the sentences occurring in a potential explanation are assumed, of course, to be empirical in the broad sense of belonging to some language adequate to the purposes of empirical science. On the problem of characterizing such systems more explicitly, see especially Scheffler's stimulating essay [46].

explanation which is involved, for example, when we ask whether a tentatively proposed but as yet untried theory would be able to explain certain puzzling empirical findings.)

Next, we say that a given potential explanation is more or less highly confirmed by a given body of evidence according as its explanans is more or less highly confirmed by the evidence in question. If the explanation is formulated in a formalized language for which an adequate quantitative concept of degree of confirmation or of inductive probability is available, we might identify the probability of the explanation relative to e with the probability of the explanans relative to e.

Finally, by a *true explanation* we understand a potential explanation with true explanans—and hence also with true explanandum.

4. *Causal Explanation and the Covering-Law Model.*

One of the various modes of explanation to which the covering-law model is relevant is the familiar procedure of accounting for an event by pointing out its "cause." In our first illustration, for example, the expansion of the lid might be said to have been caused by its immersion in hot water. Causal attributions of this sort presuppose appropriate laws, such as that whenever metal is heated under constant pressure, it expands. It is by reason of this implicit presupposition of laws that the covering-law model is relevant to the analysis of causal explanation. Let us consider this point more closely.

We will first examine general statements of causal connections, i.e., statements to the effect that an event of a given kind A—for example, motion of a magnet near a closed wire loop—will cause an event of some specified kind B—for example, flow of a current in the wire. Thereafter, we will consider statements concerning causal relations among individual events.

In the simplest case, a general statement asserting a causal connection between two kinds of events, A and B, is tantamount to the statement of the general law that whenever and wherever an instance of A occurs, it is accompanied by an instance of B. This analysis fits, for example, the statement that motion of a magnet causes a current in a neighboring wire loop. Many general statements of causal connection call for a more complex analysis, however. Thus, the statement that in a mammal, stoppage of the heart will cause death presupposes that certain "normal" conditions prevail, which are not explicitly stated, but which are surely meant to preclude, for example, the use of a heart-lung machine. "To say that X causes Y is to say that under proper conditions, an X will be followed by a Y," as Scriven[8] puts it. But unless the "proper conditions" can be specified, at least to some extent, this analysis tells us nothing about the meaning of 'X causes Y.' Now, when this kind of causal locution is used in a given context, there usually is at least some general understanding of the kind of background conditions that have to be assumed; but still, to the extent that those conditions remain indeterminate, a general statement of causal connection falls short of making a definite assertion and has at best the character of a promissory note to the effect that there are further background factors whose proper recognition would yield a truly general connection between the "cause" and "effect" under consideration.

Sentences concerning causal connections among individual events show similar characteristics. For example, the statement that the death of a certain person was caused by an overdose of phenobarbital surely presupposes a generalization, namely, a statement of a general causal connection between one kind of event, a person's taking an overdose of phenobarbital, and another, the death of that person.

Here again, the range of application for the general causal statement is not precisely stated, but a sharper specification can be given by indicating what constitutes an overdose of phenobarbital for a person—this will depend, among other things, on his weight and on his habituation to the drug—and by adding the proviso that death will result from taking such an overdose if the organism is left to itself, which implies, in particular, that no countermeasures are taken. To explain the death in question as having been caused by the antecedent taking of phenobarbital is therefore to claim that the explanandum event followed according to law upon certain antecedent circumstances. And this argument, when stated explicitly, conforms to the covering-law model.

Generally, the assertion of a causal connection between individual events seems to me unintelligible unless it is taken to make, at least implicitly, a nomological claim to the effect that there are laws which provide the basis for the causal connection asserted. When an individual event, say b, is said to have been caused by a certain antecedent event, or configuration of events, a, then surely the claim is intended that

[8] [49], p. 185.

whenever "the same cause" is realized, "the same effect" will recur. This claim cannot be taken to mean that whenever a recurs then so does b; for a and b are individual events at particular spatio-temporal locations and thus occur only once. Rather, a and b are, in this context, viewed as particular events of certain *kinds*—e.g., the expansion of a piece of metal or the death of a person—of which there may be many further instances. And the law tacitly implied by the assertion that b, as an event of kind B, was caused by a, as an event of kind A, is a general statement of causal connection to the effect that, under suitable circumstances, an instance of A is invariably accompanied by an instance of B. In most causal explanations offered in other than advanced scientific contexts, the requisite circumstances are not fully stated; for these cases, the import of the claim that b, as an instance of B, was caused by a may be suggested by the following approximate formulation: event b was in fact preceded by an event a of kind A, and by certain further circumstances which, though not fully specified or specifiable, were of such a kind that an occurrence of an event of kind A under such circumstances is universally followed by an event of kind B. For example, the statement that the burning (event of kind B) of a particular haystack was caused by a lighted cigarette carelessly dropped into the hay (particular event of kind A) asserts, first of all, that the latter event did take place; but a burning cigarette will set a haystack on fire only if certain further conditions are satisfied, which cannot at present be fully stated; and thus, the causal attribution at hand implies, second, that further conditions of a not fully specifiable kind were realized, under which an event of kind A will invariably be followed by an event of kind B.

To the extent that a statement of individual causation leaves the relevant antecedent conditions—and thus also the requisite explanatory laws—indefinite, it is like a note saying that there is a treasure hidden somewhere. Its significance and utility will increase as the location of the treasure is narrowed down, as the revelant conditions and the corresponding covering laws are made increasingly explicit. In some cases, such as that of the barbiturate poisoning, this can be done quite satisfactorily; the covering-law structure then emerges, and the statement of individual causal connection becomes amenable to test. When, on the other hand, the relevant conditions or laws remain largely indefinite, a statement of causal connection is rather in the nature of a program, or of a sketch, for an explanation in terms of causal laws; it might also be viewed as a "working hypothesis" which may prove its worth by giving new, and fruitful, direction to further research.

I would like to add here a brief comment on Scriven's observation that "when one asserts that X causes Y one is certainly committed to the generalization that an identical cause would produce an identical effect, but this in no way commits one to any necessity for producing laws not involving the term 'identical,' which justify this claim. Producing laws is one way, not necessarily more conclusive, and usually less easy than other ways of supporting the causal statement."[9] I think we have to distinguish here two questions, namely (i) what is being claimed by the statement that X causes Y, and in particular, whether asserting it commits one to a generalization, and (ii) what kind of evidence would support the causal statement, and in particular, whether such support can be provided only by producing generalizations in the form of laws.

As for the first question, I think the causal statement does imply the claim that an appropriate law or set of laws holds by virtue of which X causes Y; but, for reasons suggested above, the law or laws in question cannot be expressed by saying that an identical cause would produce an identical effect. Rather, the general claim implied by the causal statement is to the effect that there are certain "relevant" conditions of such a kind that whenever they occur in conjunction with an event of kind X, they are invariably followed by an event of kind Y.

In certain cases, some of the laws that are claimed to connect X and Y may be explicitly statable—as, for example, in our first illustration, the law that metals expand upon heating; and then, it will be possible to provide evidential support (or else disconfirmation) for them by the examination of particular instances; thus, while laws are implicitly claimed to underlie the causal connection in question, the claim can be supported by producing appropriate empirical evidence consisting of particular cases rather than of general laws. When, on the other hand, a nomological claim made by a causal statement has merely the character of an existential statement to the effect that there are relevant factors and suitable laws connecting X and Y, then it may be possible to lend some credibility to this claim by showing that under certain conditions an event of kind X is at least very frequently accompanied by an event of kind Y. This might justify the working hypothesis that the background

[9] *Ibid.*, p. 194.

conditions could be further narrowed down in a way that would eventually yield a strictly causal connection. It is this kind of statistical evidence, for example, that is adduced in support of such claims as that cigarette smoking is "a cause of" or "a causative factor in" cancer of the lung. In this case, the supposed causal laws cannot at present be explicitly stated. Thus, the nomological claim implied by this causal conjecture is of the existential type; it has the character of a working hypothesis that gives direction to further research. The statistical evidence adduced lends support to the hypothesis and justifies the program, which clearly is the aim of further research, of determining more precisely the conditions under which smoking will lead to cancer of the lung.

The most perfect examples of explanations conforming to the covering-law model are those provided by physical theories of deterministic character. A theory of this kind deals with certain specified kinds of physical systems, and limits itself to certain aspects of these, which it represents by means of suitable parameters; the values of these parameters at a given time specify the state of the system at that time; and a deterministic theory provides a system of laws which, given the state of an isolated system at one time, determine its state at any other time. In the classical mechanics of systems of mass points, for example, the state of a system at a given time is specified by the positions and momenta of the component particles at that time; and the principles of the theory—essentially the Newtonian laws of motion and of gravitation—determine the state of an isolated system of mass points at any time provided that its state at some one moment is given; in particular, the state at a specified moment may be fully explained, with the help of the theoretical principles in question, by reference to its state at some earlier time. In this theoretical scheme, the notion of a cause as a more or less narrowly circumscribed antecedent event has been replaced by that of some antecedent state of the total system, which provides the "initial conditions" for the computation, by means of the theory, of the later state that is to be explained; if the system is not isolated, i.e., if relevant outside influences act upon the system during the period of time from the initial state invoked to the state to be explained, then the particular circumstances that must be stated in the explanans include also those "outside influences"; and it is these "boundary conditions" in conjunction with the "initial" conditions which replace the everyday notion of cause, and which have to be thought of as being specified by the statements C_1,

$C_2, \ldots, C_k$ in the schematic representation (2.1) of the covering-law model.

Causal explanation in its various degrees of explicitness and precision is not the only type of explanation, however, to which the covering-law model is relevant. For example, as was noted earlier, certain empirical regularities, such as that represented by Galileo's law, can be explained by deductive subsumption under more comprehensive laws or theoretical principles; frequently, as in the case of the explanation of Kepler's laws by means of the law of gravitation and the laws of mechanics, the deduction yields a conclusion of which the generalization to be explained is only an approximation. Then the explanatory principles not only show why the presumptive general law holds, at least in approximation, but also provide an explanation for the deviations.

Another noncausal species of explanation by covering laws is illustrated by the explanation of the period of swing of a given pendulum by reference to its length and to the law that the period of a mathematical pendulum is proportional to the square root of its length. This law expresses a mathematical relation between the length and the period (a dispositional characteristic) of a pendulum *at the same time*; laws of this kind are sometimes referred to as *laws of coexistence*, in contradistinction to *laws of succession*, which concern the changes that certain systems undergo in the course of time. Boyle's, Charles's, and Van der Waals's laws for gases, which concern concurrent values of pressure, volume, and temperature of a gas; Ohm's law; and the law of Wiedemann and Franz (according to which, in metals, electric conductivity is proportional to thermal conductivity) are examples of laws of coexistence. Causal explanation in terms of antecedent events clearly calls for laws of succession in the explanans; in the case of the pendulum, where only a law of coexistence is invoked, we would not say that the pendulum's having such and such a length at a given time *caused* it to have such and such a period.[10]

It is of interest to note that in the example at hand, a statement of

[10] Note, however, that from a law of coexistence connecting certain parameters it is possible to derive laws of succession concerning the rates of change of those parameters. For example, the law expressing the period of a mathematical pendulum as a function of its length permits the derivation, by means of the calculus, of a further law to the effect that if the length of the pendulum changes in the course of time, then the rate of change of its period at any moment is proportional to the rate of change of its length, divided by the square root of its length, at that moment.

the length of a given pendulum in conjunction with the law just referred to will much more readily be accepted as explaining the pendulum's period, than a statement of the period in conjunction with the same law would be considered as explaining the length of the pendulum; and this is true even though the second argument has the same logical structure as the first: both are cases of deductive subsumption, in accordance with the schema (2.1), under a law of coexistence. The distinction made here seems to me to result from the consideration that we might change the length of the pendulum at will and thus control its period as a "dependent variable," whereas the reverse procedure does not seem possible. This idea is open to serious objections, however; for clearly, we can also change the period of a given pendulum at will, namely, by changing its length; and in doing so, we will change its length. It is not possible to retort that in the first case we have a change of length independently of a change of the period; for if the location of the pendulum, and thus the gravitational force acting on the pendulum bob, remains unchanged, then the length cannot be changed without also changing the period. In cases such as this, the common-sense conception of explanation appears to provide no clear and reasonably defensible grounds on which to decide whether a given argument that deductively subsumes an occurrence under laws is to qualify as an explanation.

The point that an argument of the form (2.1), even if its premises are assumed to be true, would not always be considered as constituting an explanation is illustrated even more clearly by the following example, which I owe to my colleague Mr. S. Bromberger. Suppose that a flagpole stands vertically on level ground and subtends an angle of 45 degrees when viewed from the ground level at a distance of 80 feet. This information, in conjunction with some elementary theorems of geometry, implies deductively that the pole is 80 feet high. The theorems in question must here be understood as belonging to physical geometry and thus as having the status of general laws, or, better, general theoretical principles, of physics. Hence, the deductive argument is of the type (2.1). And yet, we would not say that its premises *explained* the fact that the pole is 80 feet high, in the sense of showing why it is that the pole has a height of 80 feet. Depending on the context in which it is raised, the request for an explanation might call here for some kind of causal account of how it came about that the pole was given this height, or perhaps for a statement of the purpose for which this height was chosen. An account of the latter kind would again be a special case of causal explanation, invoking among the antecedent conditions certain dispositions (roughly speaking, intentions, preferences, and beliefs) on the part of the agents involved in erecting the flagpole.

The geometrical argument under consideration is not of a causal kind; in fact, it might be held that if the particular facts and the geometrical laws here invoked can be put into an explanatory connection at all, then at best we might say that the height of the pole—in conjunction with the other particulars and the laws—explains the size of the substended angle, rather than vice versa. The consideration underlying this view would be similar to that mentioned in the case of the pendulum: It might be said that by changing the height of the pole, a change in the angle can be effected, but not vice versa. But here as in the previous case, this contention is highly questionable. Suppose that the other factors involved, especially the distance from which the pole is viewed, are kept constant; then the angle can be changed, namely by changing the length of the pole; and thus, if the angle is made to change, then, trivially, the length of the pole changes. The notion that somehow we can "independently" control the length and thus make the angle a dependent variable, but not conversely, does not seem to stand up under closer scrutiny.

In sum then, we have seen that among those arguments of the form (2.1) which are not causal in character there are some which would not ordinarily be considered as even potential explanations; but ordinary usage appears to provide no clear general criterion for those arguments which are to be qualified as explanatory. This is not surprising, for our everyday conception of explanation is strongly influenced by preanalytic causal and teleological ideas; and these can hardly be expected to provide unequivocal guidance for a more general and precise analysis of scientific explanation and prediction.

5. *Covering Laws: Premises or Rules?*

Even if it be granted that causal explanations presuppose general laws, it might still be argued that many explanations of particular occurrences as formulated in everyday contexts or even in scientific discourse limit themselves to adducing certain particular facts as the presumptive causes of the explanandum event, and that therefore a formal model should construe these explanations as accounting for the explanandum by means of suitable statements of particular fact, $C_1, C_2, \ldots, C_k$, alone. Laws

would have to be cited, not in the context by *giving* such an explanation, but in the context of *justifying* it; they would serve to show that the antecedent circumstances specified in the explanans are indeed connected by causal laws with the explanandum event. Explanation would thus be comparable to proof by logical deduction, where explicit reference to the rules or laws of logic is called for, not in stating the successive steps of the proof, but only in justifying them, i.e., in showing that they conform to the principles of deductive inference. This conception would construe general laws and theoretical principles, not as scientific statements, but rather as extralogical rules of scientific inference. These rules, in conjunction with those of formal logic, would govern inferences—explanatory, predictive, retrodictive, etc.—that lead from given statements of particular fact to other statements of particular fact.

The conception of scientific laws and theories as rules of inference has been advocated by various writers in the philosophy of science.[11] In particular, it may be preferred by those who hesitate, on philosophic grounds, to accord the status of bona fide statements, which are either

[11] Among these is Schlick [48], who gives credit to Wittgenstein for the idea that a law of nature does not have the character of a statement, but rather that of an instruction for the formation of statements. Schlick's position in this article is prompted largely by the view that a genuine statement must be definitively verifiable—a condition obviously not met by general laws. But this severe verifiability condition cannot be considered as an acceptable standard for scientific statements.

More recently, Ryle—see, for example, [44], pp. 121–123—has described law statements as statements which are true or false, but one of whose jobs is to serve as inference tickets: they license their possessors to move from the assertion of some factual statements to the assertion of others.

Toulmin [53], has taken the view, more closely akin to Schlick's, that laws of nature and physical theories do not function as premises in inferences leading to observational statements, but serve as modes of representation and as rules of inference according to which statements of empirical fact may be inferred from other such statements. An illuminating discussion of this view will be found in E. Nagel's review of Toulmin's book, in *Mind*, 63:403–412 (1954); it is reprinted in Nagel [35], pp. 303–315.

Carnap [5], par. 51, makes explicit provision for the construction of languages with extralogical rules of inferences. He calls the latter physical rules, or P-rules, and emphasizes that whether, or to what extent, P-rules are to be countenanced in constructing a language is a question of expedience. For example, adoption of P-rules may oblige us to alter the rules—and thus the entire formal structure—of the language of science in order to account for some new empirical findings which, in a language without P-rules, would prompt only modification or rejection of certain statements previously accepted in scientific theory.

The admission of material rules of inference has been advocated by W. Sellars in connection with his analysis of subjunctive conditionals; see [51, 52]. A lucid general account and critical appraisal of various reasons that have been adduced in support of construing general laws as inference rules will be found in Alexander [1].

true or false, to sentences which purport to express either laws covering an infinity of potential instances or theoretical principles about unobservable "hypothetical" entities and processes.[12]

On the other hand, it is well known that in rigorous scientific studies in which laws or theories are employed to explain or predict empirical phenomena, the formulas expressing laws and theoretical principles are used, not as rules of inference, but as statements—especially as premises—quite on a par with those sentences which presumably describe particular empirical facts or events. Similarly, the formulas expressing laws also occur as conclusions in deductive arguments; for example, when the laws governing the motion of the components of a double star about their common center of gravity are derived from broader laws of mechanics and of gravitation.

It might also be noted here that a certain arbitrariness is involved in any method of drawing a line between those formulations of empirical science which are to count as statements of particular fact and those which purport to express general laws, and which accordingly are to be construed as rules of inference. For any term representing an empirical characteristic can be construed as dispositional, in which case a sentence containing it acquires the status of a generalization. Take, for example, sentences which state the boiling point of helium at atmospheric pressure, or the electric conductivity of copper: are these to be construed as empirical statements or rather as rules? The latter status could be urged on the grounds that (i) terms such as 'helium' and 'copper' are dispositional, so that their application even to one particular object involves a universal assertion, and that (ii) each of the two statements attributes a specific disposition to *any* body of helium or of copper at *any* spatio-temporal location, which again gives them the character of general statements.

The two conceptions of laws and theories—as statements or as rules of inference—correspond to two different formal reconstructions, or models, of the language of empirical science; and a model incorporating laws and theoretical principles as rules can always be replaced by one which includes them instead as scientific statements.[13] And what matters for our present purposes is simply that in either mode of representation, explanations of the kind here considered "presuppose" general theoretical principles essentially: either as indispensable premises or as indispensable rules of inference.

[12] For detailed discussions of these issues, see Barker [2], especially Ch. 7; Scheffler [46], especially Secs. 13–18; Hempel [25], especially Sec. 10.

[13] On this point, see the review by Nagel mentioned in fn. 11.

Of the two alternative construals of laws and theories, the one which gives them the status of statements seems to me simpler and more perspicuous for the analysis of the issues under investigation here; I will therefore continue to construe deductive-nomological explanations as having the form (2.1).

6. *Explanation, Prediction, Retrodiction, and Deductive Systematization—a Puzzle about 'About.'*

In a deductive-nomological explanation of a particular past event, the explanans logically implies the occurrence of the explanandum event; hence we may say of the explanatory argument that it could also have served as a predictive one in the sense that it could have been used to predict the explanandum event if the laws and particular circumstances adduced in its explanans had been taken into account at a suitable earlier time.[14] Predictive arguments of the form (2.1) will be called *deductive-nomological predictions*, and will be said to conform to the covering-law model of prediction. There are other important types of scientific prediction; among these, statistical prediction, along with statistical explanation, will be considered later.

Deductive-nomological explanation in its relation to prediction is instructively illustrated in the fourth part of the *Dialogues Concerning Two New Sciences*. Here, Galileo develops his laws for the motion of projectiles and deduces from them the corollary that if projectiles are fired from the same point with equal initial velocity, but different elevations, the maximum range will be attained when the elevation is 45°. Then, Galileo has Sagredo remark: "From accounts given by gunners, I was already aware of the fact that in the use of cannon and mortars, the maximum range . . . is obtained when the elevation is 45° . . . but to understand why this happens far outweighs the mere information obtained by the testimony of others or even by repeated experiment."[15] The reasoning that affords such understanding can readily be put into

[14] This remark does not hold, however, when all the laws invoked in the explanans are laws of coexistence (see Sec. 4) and all the particular statements adduced in the explanans pertain to events that are simultaneous with the explanandum event. I am indebted to Mr. S. Bromberger for having pointed out to me this oversight in my formulation.

[15] [18], p. 265.

the form (2.1); it amounts to a deduction, by logical and mathematical means, of the corollary from a set of premises which contains (i) the fundamental laws of Galileo's theory for the motion of projectiles and (ii) particular statements specifying that all the missiles considered are fired from the same place with the same initial velocity. Clearly then, the phenomenon previously noted by the gunners is here *explained*, and thus *understood*, by showing that its occurrence was to be expected, under the specified circumstances, in view of certain general laws set forth in Galileo's theory. And Galileo himself points with obvious pride to the predictions that may in like fashion be obtained by deduction from his laws; for the latter imply "what has perhaps never been observed in experience, namely, that of other shots those which exceed or fall short of 45° by equal amounts have equal ranges." Thus, the explanation afforded by Galileo's theory "prepares the mind to understand and ascertain other facts without need of recourse to experiment,"[16] namely, by deductive subsumption under the laws on which the explanation is based.

We noted above that if a deductive argument of the form (2.1) explains a past event, then it could have served to predict it if the information provided by the explanans had been available earlier. This remark makes a purely logical point; it does not depend on any empirical assumptions. Yet it has been argued, by Rescher, that the thesis in question "rests upon a tacit but unwarranted assumption as to the nature of the physical universe."[17]

The basic reason adduced for this contention is that "the explanation of events is oriented (in the main) towards the past, while prediction is oriented towards the future,"[18] and that, therefore, before we can decide whether (deductive-nomological) explanation and prediction have the same logical structure, we have to ascertain whether the natural laws of our world do in fact permit inferences from the present to the future as well as from the present to the past. Rescher stresses that a given system might well be governed by laws which permit deductive inferences concerning the future, but not concerning the past, or conversely; and on this point he is quite right. As a schematic illustration, consider a model "world" which consists simply of a sequence of colors, namely, Blue (B),

[16] *Ibid.*
[17] [42], p. 282.
[18] *Ibid.*, p. 286.

Green (G), Red (R), and Yellow (Y), which appear on a screen during successive one-second intervals i_1, i_2, i_3, . . . Let the succession of colors be governed by three laws:

(L_1) B is always followed by G.
(L_2) G and R are always followed by Y.
(L_3) Y is always followed by R.

Then, given the color of the screen for a certain interval, say i_3, these laws unequivocally determine the "state of the world," i.e., the screen color, for all later intervals, but not for all earlier ones. For example, given the information that during i_3 the screen is Y, the laws predict the colors for the subsequent intervals uniquely as RYRYRY . . .; but for the preceding states i_1 and i_2, they yield no unique information, since they allow here two possibilities: BG and YR.

Thus, it is possible that a set of laws governing a given system should permit unique deductive *predictions* of later states from a given one, and yet not yield unique deductive *retrodictions* concerning earlier states; conversely, a set of laws may permit unique retrodiction, but no unique prediction. But—and here lies the flaw in Rescher's argument—this is by no means the same thing as to say that such laws, while permitting deductive prediction of later states from a given one, do not permit explanation; or, in the converse case, that while permitting explanation, they do not permit prediction. To illustrate by reference to our simple model world: Suppose that during i_3 we find the screen to be Y, and that we seek to explain this fact. This can be done if we can ascertain, for example, that the color for i_1 had been B; for from the statement of this particular antecedent fact we can infer, by means of L_1, that the color for i_2 must have been G and hence, by L_2, that the color for i_3 had to be Y. Evidently, the same argument, used before i_3, could serve to predict uniquely the color for i_3 on the basis of that for i_1. Indeed, quite generally, any predictive argument made possible by the laws for our model world can also be used for explanatory purposes and vice versa. And this is so although those laws, while permitting unique predictions, do not always permit unique retrodictions. Thus, the objection under consideration misses its point because it tacitly confounds explanation with retrodiction.[19]

[19] In Sec. 3 of *SLE*, to which Rescher refers in his critique, an explanation of a past event is explicitly construed as a deductive argument inferring the occurrence of the event from "antecedent conditions" and laws; so that the temporal direction of

The notion of scientific retrodiction, however, is of interest in its own right; and, as in the case of explanation and prediction, one important variety of it is the deductive-nomological one. It has the form (2.1), but with the statements $C_1, C_2, \ldots, C_k$ referring to circumstances which occur later than the event specified in the conclusion E. In astronomy, an inference leading, by means of the laws of celestial mechanics, from data concerning the present positions and movements of the sun, the earth, and Mars to a statement of the distance between earth and Mars a year later or a year earlier illustrates deductive-nomological prediction and retrodiction, respectively; in this case, the same laws can be used for both purposes because the processes involved are reversible.

It is of interest to observe here that in their predictive and retrodictive as well as in their explanatory use, the laws of classical mechanics, or other sets of deterministic laws for physical systems, require among the premises not only a specification of the state of the system for some time, t_0, earlier or later than the time, say t_1, for which the state of the system is to be inferred, but also a statement of the boundary conditions prevailing between t_0 and t_1; these specify the external influences acting upon the system during the time in question. For certain purposes in astronomy, for example, the disturbing influence of celestial objects other than those explicitly considered may be neglected as insignificant, and the system under consideration may then be treated as "isolated"; but this should not lead us to overlook the fact that even those laws and theories of the physical sciences which provide the exemplars of deductively nomological prediction do not enable us to forecast certain future events strictly on the basis of information about the present: the predictive argument also requires certain premises concerning the future—e.g., absence of disturbing influences, such as a collision of Mars with an unexpected comet—and the temporal scope of these boundary conditions

the inference underlying explanation is the same as that of a predictive nomological argument, namely, from statements concerning certain initial (and boundary) conditions to a statement concerning the *subsequent* occurrence of the explanandum event.

I should add, however, that although all this is said unequivocally in *SLE*, there is a footnote in *SLE*, Sec. 3, which is certainly confusing, and which, though not referred to by Rescher, might have encouraged him in his misunderstanding. The footnote, numbered 2a, reads: "The logical similarity of explanation and prediction, and the fact that one is directed towards past occurrences, the other towards future ones, is well expressed in the terms 'postdictability' and 'predictability' used by Reichenbach [in *Philosophic Foundations of Quantum Mechanics*, p. 13]." To reemphasize the point at issue: postdiction, or retrodiction, is not the same thing as explanation.

must extend up to the very time at which the predicted event is to occur. The assertion therefore that laws and theories of deterministic form enable us to predict certain aspects of the future from information about the present has to be taken with a considerable grain of salt. Analogous remarks apply to deductive-nomological retrodiction and explanation.

I will use the term 'deductive-nomological systematization' to refer to any argument of the type (2.1), irrespective of the temporal relations between the particular facts specified by $C_1, C_2, \ldots, C_k$ and the particular events, if any, described by E. And, in obvious extension of the concepts introduced in Section 3 above, I will speak of *potential* (deductive-nomological) *systematizations*, of *true systematizations*, and of *systematizations* whose joint premises are more or less well confirmed by a given body of evidence.

To return now to the characterization of an explanation as a potential prediction: Scriven[20] bases one of his objections to this view on the observation that in the causal explanation of a given event (e.g., the collapse of a bridge) by reference to certain antecedent circumstances (e.g., excessive metal fatigue in one of the beams) it may well happen that the only good reasons we have for assuming that the specified circumstances were actually present lie in our knowledge that the explanandum event did take place. In this situation, we surely could not have used the explanans predictively since it was not available to us before the occurrence of the event to be predicted. This is an interesting and important point in its own right; but in regard to our conditional thesis that an explanation could have served as a prediction *if* its explanans had been taken account of in time, the argument shows only that the thesis is sometimes counterfactual (i.e., has a false antecedent), but not that it is false.

In a recent article, Scheffler[21] has subjected the idea of the structural equality of explanation and prediction to a critical scrutiny; and I would like to comment here briefly on at least some of his illuminating observations.

Scheffler points out that a prediction is usually understood to be an assertion rather than an argument. This is certainly the case; and we might add that, similarly, an explanation is often formulated, not as an argument, but as a statement, which will typically take the form 'q be-

[20] [50].
[21] [47].

cause p.' But predictive statements in empirical science are normally established by inferential procedures (which may be deductive or inductive in character) on the basis of available evidence; thus, there arises the question as to the logic of predictive arguments in analogy to the problem of the logic of explanatory arguments; and the idea of structural equality should be understood as pertaining to explanatory, predictive, retrodictive, and related arguments in science.

Scheffler also notes that a scientific prediction statement may be false, whereas, under the requirement of truth for explanations as laid down in Section 3 of *SLE*, no explanation can be false. This remark is quite correct; however, I consider it to indicate, not that there is a basic discrepancy between explanation and prediction, but that the requirement of truth for scientific explanations is unduly restrictive. The restriction is avoided by the approach that was proposed above in Section 3, and again in the present section in connection with the general characterization of scientific systematization; this approach enables us to speak of explanations no less than of predictions as being possibly false, and as being more or less well confirmed by the empirical evidence at hand.

Another critical observation Scheffler puts forth concerns the view, presented in *SLE*, that the difference between an explanatory and a predictive argument does not lie in its logical structure, but is "of a pragmatic character. If . . . we know that the phenomenon described by E has occurred, and a suitable set of statements C_1, C_2, . . ., C_k, L_1, L_2, . . ., L_r is provided afterwards, we speak of an explanation of the phenomenon in question. If the latter statements are given and E is derived prior to the occurrence of the phenomenon it describes, we speak of a prediction."[22] This characterization would make explanation and prediction mutually exclusive procedures, and Scheffler rightly suggests that they may sometimes coincide, since, for example, one may reasonably be said to be both predicting and explaining the sun's rising when, in reply to the question 'Why will the sun rise tomorrow?' one offers the appropriate astronomical information.[23]

I would be inclined to say, therefore, that in an explanation of the deductive-nomological variety, the explanandum event—which may be past, present, or future—is taken to be "given," and a set of laws and particular statements is then adduced which provides premises in an

[22] *SLE*, Sec.3

[23] Scheffler [47], p. 300.

appropriate argument of type (2.1); whereas in the case of prediction, it is the premises which are taken to be "given," and the argument then yields a conclusion about an event to occur after the presentation of the predictive inference. Retrodiction may be construed analogously. The argument referred to by Scheffler about tomorrow's sunrise may thus be regarded, first of all, as predicting the event on the basis of suitable laws and presently available information about antecedent circumstances; then, taking the predicted event as "given," the premises of the same argument constitute an explanans for it.

Thus far, I have dealt with the view that an explanatory argument is also a (potentially) predictive one. Can it be held equally that a predictive argument always offers a potential explanation? In the case of deductive-nomological predictions, an affirmative answer might be defended, though as was illustrated at the end of Section 4, there are some deductive systematizations which one would readily accept as predictions while one would find it at least awkward to qualify them as explanations. Construing the question at hand more broadly, Scheffler, and similarly Scriven,[24] have rightly pointed out, in effect, that certain sound predictive arguments of the nondeductive type cannot be regarded as affording potential explanations. For example, from suitable statistical data on past occurrences, it may be possible to "infer" quite soundly certain predictions concerning the number of male births, marriages, or traffic deaths in the United States during the next month; but none of these arguments would be regarded as affording even a low-level explanation of the occurrences they serve to predict. Now, the inferences here involved are inductive rather than deductive in character; they lead from information about observed finite samples to predictions concerning as yet unobserved samples of a given population. However, what bars them from the role of potential explanations is not their inductive character (later I will deal with certain explanatory arguments of inductive form) but the fact that they do not invoke any general laws either of strictly universal or of statistical form: it appears to be characteristic of an explanation, though not necessarily of a prediction, that it present the inferred phenomena as occurring in conformity with general laws.

In concluding this section, I would like briefly to call attention to a puzzle concerning a concept that was taken for granted in the preceding

[24] See *ibid.*, p. 296; Scriven [49].

discussion, for example, in distinguishing between prediction and retrodiction. In drawing that distinction, I referred to whether a particular given statement, the conclusion of an argument of form (2.1), was "about" occurrences at a time earlier or later than some specified time, such as the time of presentation of that argument. The meaning of this latter criterion appears at first to be reasonably clear and unproblematic. If pressed for further elucidation, one might be inclined to say, by way of a partial analysis, that if a sentence explicitly mentions a certain moment or period of time then the sentence is about something occurring at that time. It seems reasonable, therefore, to say that the sentence 'The sun rises on July 17, 1958,' says something about July 17, 1958, and that, therefore, an utterance of this sentence on July 16, 1958, constitutes a prediction.

Now the puzzle in question, which might be called the puzzle of 'about,' shows that this criterion does not even offer a partially satisfactory explication of the idea of what time a given statement is about. For example, the statement just considered can be equivalently restated in such a way that, by the proposed criterion, it is about July 15 and thus, if uttered on July 16, is about the past rather than about the future. The following rephrasing will do: 'The sun plus-two-rises on July 15,' where plus-two-rising on a given date is understood to be the same thing as rising two days after that date. By means of linguistic devices of this sort, statements about the future could be reformulated as statements about the past, or conversely; we could even replace all statements with temporal reference by statements which are, all of them, ostensibly "about" one and the same time.

The puzzle is not limited to temporal reference, but arises for spatial reference as well. For example, a statement giving the mean temperature at the North Pole can readily be restated in a form in which it speaks ostensibly about the South Pole; one way of doing this is to attribute to the South Pole the property of having, in such and such a spatial relation to it, a place where the mean temperature is such and such; another device would be to use a functor, say 'm,' which, for the South Pole, takes as its value the mean temperature at the North Pole. Even more generally there is a method which, given any particular object o, will reformulate any statement in such a way that it is ostensibly about o. If, for example, the given statement is 'The moon is spherical,' we introduce a property term, 'moon-spherical,' with the understanding that it

is to apply to o just in case the moon is spherical; the given statement then is equivalent to 'o is moon-spherical.'

The puzzle is mentioned here in order to call attention to the difficulties that face an attempt to explicate the idea of what a statement is "about," and in particular, what time it refers to; and that idea seems essential for the characterization of prediction, retrodiction, and similar concepts.[25]

Part II. Statistical Systematization

7. *Laws of Strictly General and Statistical Form.*

The nomological statements adduced in the explanans of a deductive-nomological explanation are all of a strictly general form: they purport to express strictly unexceptionable laws or theoretical principles interconnecting certain characteristics (i.e., qualitative or quantitative properties or relations) of things or events. One of the simplest forms a statement of this kind can take is that of a universal conditional: 'All (instances of) F are (instances of) G.' When the attributes in question are quantities, their interconnections are usually expressed in terms of mathematical functions, as is illustrated by many of the laws and theoretical principles of the physical sciences and of mathematical economics.

On the other hand, there are important scientific hypotheses and theoretical principles which assert that certain characters are associated, not unexceptionally or universally, but with a specified long-range frequency; we will call them statistical generalizations, or laws (or theoretical principles) of statistical form, or (statistical) probability statements. The laws of radioactive decay, the fundamental principles of quantum mechanics, and the basic laws of genetics are examples of such probability statements. These statistical generalizations, too, are used in science for the systematization of various empirical phenomena. This is illustrated, for example, by the explanatory and predictive applications of quantum theory and of the basic laws of genetics as well as by the postdictive use of the laws of radioactive decay in dating archeological relics by means of the radio-carbon method.

The rest of this essay deals with some basic problems in the logic of

[25] Professor Nelson Goodman, to whom I had mentioned my difficulties with the notion of a statement being "about" a certain subject, showed me a draft of an article entitled "About," which has now appeared in *Mind*, 70:1–24 (1961); in it, he proposes an analysis of the notion of aboutness which will no doubt prove helpful in dealing with the puzzle outlined here, and which may even entirely resolve it.

statistical systematizations, i.e., of explanatory, predictive, or similar arguments which make essential use of statistical generalizations.

Just as in the case of deductive-nomological systematization, arguments of this kind may be used to account not only for particular facts or events, but also for general regularities, which, in this case, will be of a statistical character. For example, from statistical generalizations stating that the six different results obtainable by rolling a given die are equiprobable and statistically independent of each other, it is possible to deduce the statistical generalization that the probability of rolling two aces in succession is 1/36; thus the latter statistical regularity is accounted for by subsumption (in this case purely deductive) under broader statistical hypotheses.

But the peculiar logical problems concerning statistical systematization concern the role of probability statements in the explanation, prediction, and postdiction of individual events or finite sets of such events. In preparation for a study of these problems, I shall now consider briefly the form and function of statistical generalizations.

Statistical probability hypotheses, or statistical generalizations, as understood here, bear an important resemblance to nomic statements of strictly general form: they make a universal claim, as is suggested by the term 'statistical law,' or 'law of statistical form.' Snell's law of refraction, which is of strictly general form, is not simply a descriptive report to the effect that a certain quantitative relationship has so far been found to hold, in all cases of optical refraction, between the angle of incidence and that of refraction: it asserts that that functional relationship obtains universally, in all cases of refraction, no matter when and where they occur.[26] Analogously, the statistical generalizations of genetic theory or the probability statements specifying the half lives of various radioactive substances are not just reports on the frequencies with which certain phenomena have been found to occur in some set of past instances; rather, they serve to assert certain peculiar but universal modes of connection between certain attributes of things or events.

[26] It is sometimes argued that a statement asserting such a universal connection rests, after all, only on a finite, and necessarily incomplete, body of evidence; that, therefore, it may well have exceptions which have so far gone undiscovered, and that, consequently, it should be qualified as probabilistic, too. But this argument fails to distinguish between the claim made by a given statement and the strength of its evidential support. On the latter score, all empirical statements have to count as only more or less well supported by the available evidence; but the distinction between laws of strictly universal form and those of statistical form refers to the claim made by the statements in question: roughly speaking, the former attribute a certain character to all members of a specified class; the latter, to a fixed proportion of its members.

A statistical generalization of the simplest kind asserts that the probability for an instance of F to be an instance of G is r, or briefly that $p(G,F) = r$; this is intended to express, roughly speaking, that the proportion of those instances of F which are also instances of G is r. This idea requires clarification, however, for the notion of the proportion of the (instances of) G among the (instances of) F has no clear meaning when the instances of F do not form a finite class. And it is characteristic of probability hypotheses with their universal character, as distinguished from statements of relative frequencies in some finite set, that the reference class—F in this case—is not assumed to be finite; in fact, we might underscore their peculiar character by saying that the probability r does not refer to the class of all actual instances of F but, so to speak, to the class of all its potential instances.

Suppose, for example, that we are given a homogeneous regular tetrahedron whose faces are marked 'I,' 'II,' 'III,' 'IV.' We might then be willing to assert that the probability of obtaining a III, i.e., of the tetrahedran's coming to rest on that face, upon tossing it out of a dice box is ¼; but while this assertion would be meant to say something about the frequency with which a III is obtained as a result of rolling the tetrahedron, it could not be construed as simply specifying that frequency for the class of all tosses which are in fact ever performed with the tetrahedron. For we might well maintain our probability hypothesis even if the given tetrahedron were tossed only a few times throughout its existence, and in this case, our probability statement would certainly not be meant to imply that exactly or even nearly, one fourth of those tosses yielded the result III. In fact, we might clearly maintain the probability statement even if the tetrahedron happened to be destroyed without ever having been tossed at all. We might say, then, that the probability hypothesis ascribes to the tetrahedron a certain disposition, namely, that of yielding a III in about one out of four cases in the long run. That disposition may also be described by a subjunctive or counterfactual statement: If the tetrahedron were to be tossed (or had been tossed) a large number of times, it would yield (would have yielded) the result III in about one fourth of the cases.[27]

[27] The characterization given here of the concept of statistical probability seems to me to be in agreement with the general tenor of the "propensity interpretation" advo-

Let us recall here in passing that nomological statements of strictly general form, too, are closely related to corresponding subjunctive and counterfactual statements. For example, the lawlike statement 'All pieces of copper expand when heated' implies the subjunctive conditional 'If this copper key were heated it would expand' and the counterfactual statement, referring to a copper key that was kept at constant temperature during the past hour, 'If this copper key had been heated half an hour ago, it would have expanded.'[28]

To obtain a more precise account of the form and function of probability statements, I will examine briefly the elaboration of the concept of statistical probability in contemporary mathematical theory. This examination will lead to the conclusion that the logic of statistical systematization differs fundamentally from that of deductive-nomological systematization. One striking symptom of the difference is what will be called here *the ambiguity of statistical systematization*.

In Section 8, I will describe and illustrate this ambiguity in a general manner that presupposes no special theory of probability; then in Section 9, I will show how it reappears in the explanatory and predictive

cated by Popper in recent years. This interpretation "differs from the purely statistical or frequency interpretation only in this—that it considers the probability as a characteristic property of the experimental arrangement rather than as a property of a sequence"; the property in question is explicitly construed as *dispositional*. (Popper [39], pp. 67–68. See also the discussion of this paper at the Ninth Symposium of the Colston Research Society, in Körner [30], pp. 78–89 *passim*.) However, the currently available statements of the propensity interpretation are all rather brief (for further references, see Popper [40]); a fuller presentation is to be given in a forthcoming book by Popper.

[28] In fact, Goodman [20], has argued very plausibly that one symptomatic difference between lawlike and nonlawlike generalizations is precisely that the former are able to lend support to corresponding subjunctive or counterfactual conditionals; thus the statement 'If this copper key were to be heated it would expand' can be supported by the law mentioned above. By contrast, the general statement 'All objects ever placed on this table weigh less than one pound' is nonlawlike, i.e., even if true, it does not count as a law. And indeed, even if we knew it to be true, we would not adduce it in support of corresponding counterfactuals; we would not say, for example, that if a volume of Merriam-Webster's Unabridged Dictionary had been placed on the table, it would have weighed less than a pound. Similarly, it might be added, general statements of this latter kind possess no explanatory power: this is why the sentences L_1 L_2, . . ., L_n in the explanans of any deductive-nomological explanation are required to be lawlike.

The preceding considerations suggest the question whether there is a category of statistical probability statements whose status is comparable to that of accidental generalizations. It would seem clear, however, that insofar as statistical probability statements are construed as dispositional in the sense suggested above, they have to be considered as being analogous to lawlike statements.

use of probability hypotheses as characterized by the mathematical theory of statistical probability.

8. *The Ambiguity of Statistical Systematization.*

Consider the following argument which represents, in a nutshell, an attempt at a statistical explanation of a particular event: "John Jones was almost certain to recover quickly from his streptococcus infection, for he was given penicillin, and almost all cases of streptococcus infection clear up quickly upon administration of penicillin." The second statement in the explanans is evidently a statistical generalization, and while the probability value is not specified numerically, the words 'almost all cases' indicate that it is very high.

At first glance, this argument appears to bear a close resemblance to deductive-nomological explanations of the simplest form, such as the following: This crystal of rock salt, when put into a Bunsen flame, turns the flame yellow, for it is a sodium salt, and all sodium salts impart a yellow color to a Bunsen flame. This argument is basically of the form:

$$\text{(8.1)}\qquad \begin{array}{l} \text{All F are G.} \\ \text{x is F.} \\ \hline \text{x is G.} \end{array}$$

The form of the statistical explanation, on the other hand, appears to be expressible as follows:

$$\text{(8.2)}\qquad \begin{array}{c} \text{Almost all F are G.} \\ \text{x is F.} \\ \hline \text{x is almost certain to be G.} \end{array}$$

Despite this appearance of similarity, however, there is a fundamental difference between these two kinds of argument: A nomological explanation of the type (8.1) accounts for the fact that x is G by stating that x has another character, F, which is uniformly accompanied by G, in virtue of a general law. If in a given case these explanatory assumptions are in fact true, then it follows logically that x must be G; hence x cannot possibly possess a character, say H, in whose presence G is uniformly absent; for otherwise, x would have to be both G and non-G. In the argument (8.2), on the other hand, x is said to be almost certain to have G because it has a character, F, which is accompanied by G in almost all instances. But even if in a given case the explanatory statements are both true, x may possess, in addition to F, some other attribute, say H, which is almost always accompanied by non-G. But by the very logic underlying (8.2), this attribute would make it almost certain that x is not G.

Suppose, for example, that almost all, but not quite all, penicillin-treated, streptococcal infections result in quick recovery, or briefly, that almost all P are R; and suppose also that the particular case of illness of patient John Jones which is under discussion—let us call it j—is an instance of P. Our original statistical explanation may then be expressed in the following manner, which exhibits the form (8.2):

(8.3a)
$$\frac{\begin{array}{c}\text{Almost all P are R.}\\ \text{j is P.}\end{array}}{\text{j is almost certain to be R.}}$$

Next, let us say that an event has the property P* if it is either the event j itself or one of those infrequent cases of penicillin-treated streptococcal infection which do not result in quick recovery. Then clearly j is P*, whether or not j is one of the cases resulting in recovery, i.e., whether or not j is R. Furthermore, almost every instance of P* is an instance of non-R (the only possible exception being j itself). Hence, the argument (8.3a) in which, on our assumption, the premises are true can be matched with another one whose premises are equally true, but which by the very logic underlying (8.3a), leads to a conclusion that appears to contradict that of (8.3a):

(8.3b)
$$\frac{\begin{array}{c}\text{Almost all P* are non-R.}\\ \text{j is P*.}\end{array}}{\text{j is almost certain to be non-R.}}$$

If it should be objected that the property P* is a highly artificial property and that, in particular, an explanatory statistical law should not involve essential reference to particular individuals (such as j in our case), then another illustration can be given which leads to the same result and meets the contemplated requirement. For this purpose, consider a number of characteristics of John Jones at the onset of his illness, such as his age, height, weight, blood pressure, temperature, basal metabolic rate, and IQ. These can be specified in terms of numbers; let n_1, n_2, n_3, . . . be the specific numerical values in question. We will say that an event has the property S if it is a case of streptococcal infection in a patient

who at the onset of his illness has the height n_1, age n_2, weight n_3, blood pressure n_4, and so forth. Clearly, this definition of S in terms of numerical characteristics no longer makes reference to j. Finally, let us say that an event has the property P** if it is either an instance of S or one of those infrequent cases of streptococcal infection treated with penicillin which do not result in quick recovery. Then evidently j is P** because j is S; and furthermore, since S is a very rare characteristic, almost every instance of P** is an instance of non-R. Hence, (8.3a) can be matched with the following argument, in which the explanatory probability hypothesis involves no essential reference to particular cases:

(8.3c)
$$\frac{\begin{array}{c}\text{Almost all P** are non-R.}\\ \text{j is P**.}\end{array}}{\text{j is almost certain to be non-R.}}$$

The premises of this argument are true if those of (8.3a) are, and the conclusion again appears to be incompatible with that of (8.3a).

The peculiar phenomenon here illustrated will be called the *ambiguity of statistical explanation*. Briefly, it consists in the fact that if the explanatory use of a statistical generalization is construed in the manner of (8.2), then a statistical explanation of a particular event can, in general, be matched by another one, equally of the form (8.2), with equally true premises, which statistically explains the nonoccurrence of the same event. The same difficulty arises, of course, when statistical arguments of the type (8.2) are used for predictive purposes. Thus, in the case of our illustration, we might use either of the two arguments (8.3a) and (8.3c) in an attempt to predict the effect of penicillin treatment in a fresh case, j, of streptococcal infection; and even though both followed the same logical pattern—that exhibited in (8.2)—and both had true premises, one argument would yield a favorable, the other an unfavorable forecast. We will, therefore, also speak of the *ambiguity of statistical prediction* and, more inclusively, of the *ambiguity of statistical systematization*.

This difficulty is entirely absent in nomological systematization, as we noted above; and it evidently throws into doubt the explanatory and predictive relevance of statistical generalizations for particular occurrences. Yet there can be no question that statistical generalizations are widely invoked for explanatory and predictive purposes in such diverse fields as physics, genetics, and sociology. It will be necessary, therefore, to examine more carefully the logic of the arguments involved and, in particular, to reconsider the adequacy of the analysis suggested in (8.2). And while for a general characterization of the ambiguity of statistical explanation it was sufficient to use an illustration of statistical generalization of the vague form 'Almost all F are G,' we must now consider the explanatory and predictive use of statistical generalizations in the precise form of quantitative probability statements: 'The probability for an F to be a G is r.' This brings us to the question of the theoretical status of the statistical concept of probability.

9. *The Theoretical Concept of Statistical Probability and the Problem of Ambiguity.*

The mathematical theory of statistical probability[29] seeks to give a theoretical systematization of the statistical aspects of random experiments. Roughly speaking, a random experiment is a repeatable process which yields in each case a particular finite or infinite set of "results," in such a way that while the results vary from repetition to repetition in an irregular and practically unpredictable manner, the relative frequencies with which the different results occur tend to become more or less constant for large numbers of repetitions. The theory of probability is intended to provide a "mathematical model," in the form of a deductive system, for the properties and interrelations of such long-run frequencies, the latter being represented in the model by probabilities.

In the mathematical theory of probability, each of the different outcomes of a random experiment which have probabilities assigned to them is represented by a set of what might be called elementary possibilities. For example, if the experiment is that of rolling a die, then getting an ace, a deuce, and so forth, would normally be chosen as elementary possibilities; let us refer to them briefly as I, II, . . ., VI, and let F be the set of these six elements. Then any of those results of rolling a die to which probabilities are usually assigned can be represented by a subset of F: getting an even number, by the set (II, IV, VI); getting a prime number, by the set (II, III, V); rolling an ace, by the unit set (I); and so forth. Generally, a random experiment is represented in the theory by a set F and a certain set, F*, of its subsets, which represent the possible outcomes that have definite probabilities assigned to them. F* will sometimes, but not always, contain all the subsets of F. The mathematical theory also requires F* to contain, for each of its member sets, its complement in F; and also for any two of its member sets, say G_1 and G_2, their sum, $G_1 \vee G_2$, and their products, $G_1 \cdot G_2$. As a consequence, F* contains F as a member set.[30] The probabilities associated with the different outcomes of a random experiment then are represented by a real-valued function $p_F(G)$ which ranges over the sets in F*.

The postulates of the theory specify that p_F is a nonnegative additive set function such that $p_F(F) = 1$; i.e., for all G in F*, $p_F(G) \geqq 0$; if G_1 and G_2 are mutually exclusive sets in F* then $p_F(G_1 \vee G_2) = p_F(G_1) + p_F(G_2)$. These stipulations permit the proof of the theorems of elementary probability theory; to deal with experiments that permit infinitely many different outcomes, the requirement of additivity is suitably extended to infinite sequences of mutually exclusive member sets of F*.

The abstract theory is made applicable to empirical subject matter by means of an interpretation which connects probability statements with sentences about long-run relative frequencies associated with random experiments. I will state the interpretation in a form which is essentially that given by Cramér,[31] whose book *Mathematical Methods of Statistics* includes a detailed discussion of the foundations of mathematical probability theory and its applications. For convenience, the notation '$p_F(G)$' for the probability of G relative to F will now be replaced by 'p(G, F).'

(9.1) *Frequency interpretation of statistical probability*: Let F be a given kind of random experiment and G a possible result of it; then the statement that $p(G, F) = r$ means that in a long series

[29] The mathematical theory of statistical probability has been developed in two major forms. One of these is based on an explicit definition of probabilities as limits of relative frequencies in certain infinite reference sequences. The use of this limit definition is an ingenious attempt to permit the development of a simple and elegant theory of probability by means of the apparatus of mathematical analysis, and to reflect at the same time the intended statistical application of the abstract theory. The second approach, which offers certain theoretical advantages and is now almost generally adopted, develops the formal theory of probability as an abstract theory of certain set-functions and then specifies rules for its application to empirical subject matter. The brief characterization of the theory of statistical probability given in this section follows the second approach. However, the problem posed by the ambiguity of statistical systematization arises as well when the limit definition of probability is adopted.

[30] See, for example, Kolmogoroff [31], Sec 2.

[31] [11], pp. 148–149. Similar formulations have been given by other representatives of this measure-theoretical conception of statistical probability, for example, by Kolmogoroff [31], p. 4.

of repetitions of F, it is practically certain that the relative frequency of the result G will be approximately equal to r.

Evidently, this interpretation does not offer a precise definition of probability in statistical terms: the vague phrases 'a long series,' 'practically certain,' and 'approximately equal' preclude that. But those phrases are chosen deliberately to enable formulas stating precisely fixed numerical probability values to function as theoretical representations of near-constant relative frequencies of certain results in extended repetitions of a random experiment.

Cramér also formulates two corollaries of the above rule of interpretation; they refer to those cases where r differs very little from 0 or from 1. These corollaries will be of special interest for an examination of the question of ambiguity in the explanatory and predictive use of probability statements, and I will therefore note them here (in a form very similar to that chosen by Cramér):

(9.2a) If $0 \leq p(G, F) < \epsilon$, where ϵ is some very small number, then, if a random experiment of kind F is performed one single time, it can be considered as practically certain that the result G will not occur.[32]

(9.2b) If $1 - \epsilon < p(G, F) \leq 1$, where ϵ is some very small number, then if a random experiment of kind F is performed one single time, it can be considered as practically certain that the result G will occur.[33]

I now turn to the explanatory use of probability statements. Consider the experiment, D, of drawing, with subsequent replacement and thorough mixing, a ball from an urn containing one white ball and 99 black ones of the same size and material. Let us suppose that the probability, p(W, D), of obtaining a white ball as a result of a performance of D is .99. According to the statistical interpretation, this is an empirical hypothesis susceptible of test by reference to finite statistical samples, but for the moment, we need not enter into the question how the given hypothesis might be established. Now, rule (9.2b) would seem to indicate that this hypothesis might be used in statistically explaining or predicting the results of certain individual drawings from the urn. Suppose, for example, that a particular drawing, d, produces a white ball. Since p(W, D) differs from 1 by less than, say, .015, which is a rather small

[32] Cf. Cramer [11], p. 149; see also the very similar formulation in Kolmogoroff [31], p. 4.

[33] Cf. Cramér [11], p. 150.

number, (9.2b) yields the following argument, which we might be inclined to consider as a statistical explanation of the fact that d is W:

(9.3)

$1 - .015 < p(W, D) \leq 1$; and .015 is a very small number.

d is an instance of D.

It is practically certain that d is W.

This type of reasoning is closely reminiscent of our earlier argument (8.3a), and it leads into a similar difficulty, as will now be shown. Suppose that besides the urn just referred to, which we will assume to be marked '1,' there are 999 additional urns of the same kind, each containing 100 balls, all of which are black. Let these urns be marked '2,' '3' . . . '1000.' Consider now the experiment E which consists in first drawing a ticket from a bag containing 1000 tickets of equal size, shape, etc., bearing the numerals '1,' '2' . . . '1000,' and then drawing a ball from the urn marked with the same numeral as the ticket drawn. In accordance with standard theoretical considerations, we will assume that p(W, E) = .00099. (This hypothesis again is capable of confirmation by statistical test in view of the interpretation (9.1).) Now, let e be a particular performance of E in which the first step happens to yield the ticket numbered 1. Then, since e is an instance of E, the interpretative rule (9.2a) permits the following argument:

(9.4a)

$0 \leq p(W, E) < .001$; and .001 is a very small number.

e is an instance of E.

It is practically certain that e is not W.

But on our assumption, the event e also happens to be an instance of the experiment D of drawing a ball from the first urn; we may therefore apply to it the following argument:

(9.4b)

$1 - .015 < p(W, D) \leq 1$; and .015 is a very small number.

e is an instance of D.

It is practically certain that e is W.

Thus, in certain cases the interpretative rules (9.2a) and (9.2b) yield arguments which again exhibit what was called above the ambiguity of statistical systematization.

This ambiguity clearly springs from the fact that (a) the probability of obtaining an occurrence of some specified kind G depends on the random experiment whose result G is being considered, and that (b) a particular instance of G can normally be construed as an outcome of different kinds of random experiment, with different probabilities for the outcome in question; as a result, under the frequency interpretation given in (9.2a) and (9.2b), an occurrence of G in a particular given case may be shown to be both practically certain and practically impossible. This ambiguity does not represent a flaw in the formal theory of probability: it arises only when the empirical interpretation of that theory is brought into play.

It might be suspected that the trouble arises only when an attempt is made to apply probability statements to individual events, such as one particular drawing in our illustration: statistical probabilities, it might be held, have significance only for reasonably large samples. But surely this is unconvincing since there is only a difference in degree between a sample consisting of just one case and a sample consisting of many cases. And indeed, the problem of ambiguity recurs when probability statements are used to account for the frequency with which a specified kind G of result occurs in finite samples, no matter how large.

For example, let the probability of obtaining recovery (R) as the result of the "random experiment" P of treating cases of streptococcus infection with penicillin be $p(R, P) = .75$. Then, assuming statistical independence of the individual cases, the frequency interpretation yields the following consequence, which refers to more or less extensive samples: For any positive deviation d, however small, there exists a specifiable sample size n_d such that it is practically certain that in one single series of n_d repetitions of the experiment P, the proportion of cases of R will deviate from .75 by less than d.[34] It would seem therefore that a recovery rate of close to 75 per cent in a sufficiently large number of instances of P could be statistically explained or predicted by means of the probability statement that $p(R, P) = .75$. But any such series of instances can also be construed as a set of cases of another random experiment for which it is practically certain that almost all the cases in the sample recover; alternatively, the given cases can be construed as a set of instances of yet another random experiment for which it is practically certain that none of the cases in a sample of the given size will recover. The arguments leading to this conclusion are basically similar to those

[34] *Ibid.*, pp. 197–198.

presented in connection with the preceding illustrations of ambiguity; the details will therefore be omitted.

In its essentials, the ambiguity of statistical systematization can be characterized as follows: If a given object or set of objects has an attribute A which with high statistical probability is associated with another attribute C, then the same object or set of objects will, in general, also have an attribute B which, with high statistical probability, is associated with non-C. Hence, if the occurrence of A in the particular given case, together with the probability statement which links A with C, is regarded as constituting adequate grounds for the predictive or explanatory conclusion that C will almost certainly occur in the given case, then there exists, apart from trivial exceptions, always a competing argument which in the same manner, from equally true premises, leads to the predictive or explanatory conclusion that C will not occur in that same case. This peculiarity has no counterpart in nomological explanation: If an object or set of objects has a character A which is invariably accompanied by C then it cannot have a character B which is invariably accompanied by non-C.[35]

The ambiguity of statistical explanation should not, of course, be taken to indicate that statistical probability hypotheses have no explanatory or predictive significance, but rather that the above analysis of the logic of statistical systematization is inadequate. That analysis was suggested by a seemingly plausible analogy between the systematizing use of statistical generalizations and that of nomic ones—an analogy which seems to receive strong support from the interpretation of statistical generalizations which is offered in current statistical theory. Nevertheless, that analogy is deceptive, as will now be shown.

10. *The Inductive Character of Statistical Systematization and the Requirement of Total Evidence.*

It is typical of the statistical systematizations considered in this study that their "conclusion" begins with the phrase 'It is almost certain that,'

[35] My manuscript here originally contained the phrase 'is invariably (or even in some cases) accompanied by non-C.' By reading the critique of this passage as given in the manuscript of Professor Scriven's contribution to the present volume, I became aware that the claim made in parentheses is indeed incorrect. Since the point is entirely inessential to my argument, I deleted the parenthetical remark after having secured Professor Scriven's concurrence. However, Professor Scriven informed me that he would not have time to remove whatever references to this lapse his manuscript might contain: I therefore add this note for clarification.

which never occurs in the conclusion of a nomological explanation or prediction. The two schemata (8.1) and (8.2) above exhibit this difference in its simplest form. A nomological systematization of the form (8.1) is a deductively valid argument: if its premises are true then so is its conclusion. For arguments of the form (8.2), this is evidently not the case. Could the two types of argument be assimilated more closely to each other by giving the conclusion of (8.1) the form 'It is certain that x is G'? This suggestion involves a misconception which is one of the roots of the puzzle presented by the ambiguity of statistical systematization. For what the statement 'It is certain that x is G' expresses here can be restated by saying that the conclusion of an argument of form (8.1) cannot be false if the premises are true, i.e., that the conclusion is a logical consequence of the premises. Hence, the certainty here in question represents not a property of the conclusion that x is G, but rather a relation which that conclusion bears to the premises of (8.1). Generally, a sentence is certain, in this sense, relative to some class of sentences just in case it is a logical consequence of the latter. The contemplated reformulation of the conclusion of (8.1) would therefore be an elliptic way of saying that

(10.1) 'x is G' is certain relative to, i.e., is a logical consequence of, the two sentences 'All F are G' and 'x is F.'[36]

But clearly this is not equivalent to the original conclusion of (8.1); rather, it is another way of stating that the entire schema (8.1) is a deductively valid form of inference.

Now, the basic error in the formulation of (8.2) is clear: near certainty, like certainty, must be understood here not as a property but as a relation; thus, the "conclusion" of (8.2) is not a complete statement but an elliptical formulation of what might be more adequately expressed as follows:

(10.2) 'x is G' is almost certain relative to the two sentences 'Almost all F are G' and 'x is F.'

The near certainty here invoked is sometimes referred to as (high)

[36] A sentence of the form 'It is certain that x is G' ostensibly attributes the modality of certainty to the proposition expressed by the conclusion in relation to the propositions expressed by the premises. For the purposes of the present study, involvement with propositions can be avoided by construing the given modal sentence as expressing a logical relation that the conclusion, taken as a sentence, bears to the premise sentences. Concepts such as near certainty and probability can, and will here, equally be treated as applying to pairs of sentences rather than to pairs of propositions.

probability; the conclusion of arguments like (8.2) is then expressed by such phrases as '(very) probably, x is G,' or 'it is (highly) probable that x is G'; a nonelliptic restatement would then be given by saying that the sentences 'Almost all F are G' and 'x is F' taken jointly lend strong support to, or confer a high probability or a high degree of rational credibility upon, 'x is G.' The probabilities referred to here are logical or inductive probabilities, in contradistinction to the statistical probabilities mentioned in the premises of the statistical systematization under examination. The notion of logical probability will be discussed more fully a little later in the present section.

As soon as it is realized that the ostensible "conclusions" of arguments such as (8.2) and their quantitative counterparts, such as (9.3), are elliptic formulations of relational statements, one puzzling aspect of the ambiguity of statistical systematization vanishes: the apparently conflicting claims of matched argument pairs such as (8.3a) and (8.3b) or (9.4a) and (9.4b) do not conflict at all. For what the matched arguments in a pair claim is only that each of two contradictory sentences, such as 'j is R' and 'j is not R' in the pair (8.3), is strongly supported by certain other statements, which, however, are quite different for the first and for the second sentence in question. Thus far then, no more of a "conflict" is established by a pair of matched statistical systematizations than, say, by the following pair of deductive arguments, which show that each of two contradictory sentences is even conclusively supported, or made certain, by other suitable statements which, however, are quite different for the first and for the second sentence in question:

$$(10.3a)\quad \frac{\begin{array}{c}\text{All F are G.}\\ \text{a is F.}\end{array}}{\text{a is G.}}$$

$$(10.3b)\quad \frac{\begin{array}{c}\text{No H is G.}\\ \text{a is H.}\end{array}}{\text{a is not G.}}$$

The misconception thus dispelled arises from a misguided attempt to construe arguments containing probability statements among their premises in analogy to deductive arguments such as (8.1)—an attempt which prompts the construal of formulations such as 'j is almost certain to be

R' or 'probably, j is R' as self-contained complete statements rather than as elliptically formulated statements of a relational character.[37]

The idea, repeatedly invoked in the preceding discussion, of a statement or set of statements e (the evidence) providing strong grounds for asserting a certain statement h (the hypothesis), or of e lending strong support to h, or making h nearly certain is, of course, the central concept of the theory of inductive inference. It might be conceived in purely qualitative fashion as a relation S which h bears to e if e lends strong support to h; or it may be construed in quantitative terms, as a relation capable of gradations which represents the extent to which h is supported by e. Some recent theories of inductive inference have aimed at developing rigorous quantitative conceptions of inductive support: this

[37] These remarks seem to me to be relevant, for example, to C. I. Lewis's notion of categorical, as contradistinguished from hypothetical, probability statements. For in [32], p. 319, Lewis argues as follows: "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."

This position seems to me to be open to just those objections which have been 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 is 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 which rules out the conception of categorical probability statements.

is true especially of the systems of inductive logic constructed by Keynes and others and recently, in a particularly impressive form, by Carnap.[38] If—as in Carnap's system—the concept is construed so as to possess the formal characteristics of a probability, it will be referred to as the logical (or inductive) probability, or as the degree of confirmation, c(h, e), of h relative to e. (This inductive probability, which is a function of statements, must be sharply distinguished from statistical probability, which is a function of classes of events.) As a general phrase referring to a quantitative notion of inductive support, but not tied to any one particular theory of inductive support or confirmation, let us use the expression '(degree of) inductive support of h relative to e.'[39]

An explanation, prediction, or retrodiction of a particular event or set of events by means of principles which include statistical generalizations has then to be conceived as an inductive argument. I will accordingly speak of *inductive systematization* (in contradistinction to *deductive systematization*, where whatever is explained, predicted, or retrodicted is a deductive consequence of the premises adduced in the argument).

When it is understood that a statistical systematization is an inductive argument, and that the high probability or near certainty mentioned in the conclusions of such arguments as (8.3a) and (8.3b) is relative to the premises, then, as shown, one puzzle raised by the ambiguity of statistical explanation is resolved, namely the impression of a conflict, indeed a near incompatibility, of the claims of two equally sound inductive systematizations.

But the same ambiguity raises another, more serious, problem, which now calls for consideration. It is very well to point out that in (8.3a) and (8.3b) the contradictory statements 'j is R' and 'j is not R' are shown to be almost certain by referring to different sets of "premises": it still remains the case that both of these sets are true. Here, the analogy to (10.3a) and (10.3b) breaks down: in these deductive arguments with contradictory conclusions the two sets of premises cannot both be true. Thus, it would seem that by statistical systematizations based on suitably

[38] See especially [7, 8], and, for a very useful survey [6].

[39] In a recent study, Kemeny and Oppenheim [29], have proposed, and theoretically developed, an interesting concept of "degree of factual support" (of a hypothesis by given evidence), which differs from Carnap's concept of degree of confirmation, or inductive probability, in important respects; for example, it does not have the formal character of a probability function. For a suggestive distinction and comparison of different concepts of evidence, see Rescher [43].

chosen bodies of true information, we may lend equally strong support to two assertions which are incompatible with each other. But then—and this is the new problem—which of such alternative bodies of evidence is to be relied on for the purposes of statistical explanation or prediction?

An answer is suggested by a principle which Carnap calls *the requirement of total evidence*. It lays down a general maxim for all applications of inductive reasoning, 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."[40] Instead of the total evidence, a smaller body, e_1, of evidence may be used on condition that the remaining part, e_2, of the total evidence is inductively irrelevant to the hypothesis h whose confirmation is to be determined. If, as in Carnap's system, the degree of confirmation is construed as an inductive probability, the irrelevance of e_2 for h relative to e_1 can be expressed by the condition that $c(h, e_1 \cdot e_2) = c(h, e_1)$.[41]

The general consideration underlying the requirement of total evidence is obviously this: If an investigator wishes to decide what credence to give to an empirical hypothesis or to what extent to rely on it in planning his actions, then rationality demands that he take into account all the relevant evidence available to him; if he were to consider only part of that evidence, he might arrive at a much more favorable, or a much less favorable, appraisal, but it would surely not be rational for him to base his decision on evidence he knew to be selectively biased. In terms of the concept of degree of confirmation, the point might be stated by saying that the degree of confirmation assigned to a hypothesis by the principles of inductive

[40] Carnap [7], p. 211. In his comments, pp. 211–213, Carnap points out that in less explicit form, the requirement of total evidence has been recognized by various authors at least since Bernoulli. The idea also is suggested in the passage from Lewis [32], quoted in fn. 36. Similarly, Williams, whose book *The Ground of Induction* centers about various arguments that have the character of statistical systematizations, 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." (Williams [55], p. 72.)

I wish to acknowledge here my indebtedness to Professor Carnap, to whom I turned in 1945, when I first noticed the ambiguity of statistical explanation, and who promptly pointed out to me in a letter that this was but one of several apparent paradoxes of inductive logic which result from violations of the requirement of total evidence.

In his recent book, Barker [2], pp. 70–78, concisely and lucidly presents the gist of the puzzle under consideration here and examines the relevance to it of the principle of total evidence.

[41] Cf. Carnap [7], pp. 211, 494.

logic will represent the rational credibility of the hypothesis for a given investigator only if the argument takes into account all the relevant evidence available to the investigator.

The requirement of total evidence is not a principle of inductive logic, which is concerned with relations of potential evidential support among statements, i.e., with whether, or to what degree, a given set of statements supports a given hypothesis. Rather, the requirement is a maxim for the *application* of inductive logic; it might be said to state a necessary condition of rationality in forming beliefs and making decisions on the basis of available evidence. The requirement is not limited to arguments of the particular form of statistical systematizations, where the evidence, represented by the "premises," includes statistical generalizations: it is a necessary condition of rationality in the application of any mode of inductive reasoning, including, for example, those cases in which the evidence contains no generalizations, statistical or universal, but only data on particular occurrences.

Let me note here that in the case of deductive systematization, the requirement is automatically satisfied and thus presents no special problem.[42] For in a deductively valid argument whose premises constitute only part of the total evidence available at the time, that part provides conclusive grounds for asserting the conclusion; and the balance of the total evidence is irrelevant to the conclusion in the strict sense that if it were added to the premises, the resulting premises would still constitute conclusive grounds for the conclusion. To state this in the language of inductive logic: the logical probability of the conclusion relative to the premises of a deductive systematization is 1, and it remains 1 no matter what other parts of the total evidence may be added to the premises.

The residual problem raised by the ambiguity of probabilistic explanation can now be resolved by requiring that if a statistical systematization is to qualify as a rationally acceptable explanation or prediction (and not just as a formally sound *potential* explanation or prediction), it must satisfy the requirement of total evidence. For under this requirement, the "premises" of an acceptable statistical systematization whose "conclusion" is a hypothesis h must consist either of the total evidence e or of some subset of it which confers on h the same inductive probability as e; and

[42] Carnap [7], p. 211, says "There is no analogue to this requirement [of total evidence] in deductive logic"; but it seems more accurate to say that the requirement is automatically met here.

the same condition applies to an acceptable systematization which has the negation of h as its "conclusion." But one and the same evidence, e, cannot—if it is logically self-consistent—confer a high probability on h as well as on its negation, since the sum of the two probabilities is unity. Hence, of two statistical systematizations whose premises confer high probabilities on h and on the negation of h, respectively, at least one violates the requirement of total evidence and is thus ruled out as unacceptable.

The preceding considerations suggest that a statistical systematization may be construed generally as an inductive argument showing that a certain statement or finite set of statements, e, which includes at least one statistical law, gives strong but not logically conclusive support to a statement h, which expresses whatever is being explained, predicted, retrodicted, etc. And if an argument of this kind is to be acceptable in science as an empirically sound explanation, prediction, or the like—rather than only a formally adequate, or potential one—then it will also have to meet the requirement of total evidence.

But an attempt to apply the requirement of total evidence to statistical systematizations of the simple kind considered so far encounters a serious obstacle. This was noted, among others, by S. Barker with special reference to "statistical syllogisms," which are inductive arguments with two premises, very similar in character to the arguments (9.4a) and (9.4b) above. Barker points out, in effect, that the statistical syllogism is subject to what has been called here the ambiguity of statistical systematization, and he goes on to argue that the principle of total evidence will be of no avail as a way to circumvent this shortcoming because generally our total evidence will consist of far more than just two statements, which would moreover have to be of the particular form required for the premises of a statistical syllogism.[43] This observation would not raise a serious difficulty, at least theoretically speaking, if an appropriate general system of inductive logic were available: the rules of this system might enable us to show that that part of our total evidence which goes beyond the premises of our simple statistical argument is inductively irrelevant to the conclusion in the sense specified earlier in this section. Since no inductive logic of the requisite scope is presently at hand, however, it is a question of great interest whether some more manageable substitute for the requirement of total evidence might not be formulated which would not presuppose a full system of inductive logic and would be applicable to simple statistical systematizations. This question will be examined in the next section on the basis of a closer analysis of simple statistical systematizations offered by empirical science.

[43] See Barker [2], pp. 76–78. The point is made in a more general form by Carnap [7], p. 404.

11. *The Logical Form of Simple Statistical Systematizations: A Rough Criterion of Evidential Adequacy.*

Let us note, first of all, that empirical science offers many statistical systematizations which accord quite well with the general characterization to which we were led in the preceding section.

For example, by means of Mendelian genetic principles it can be shown that in a random sample taken from a population of pea plants each of whose parent plants represents a cross of a pure white-flowered and a pure red-flowered strain, approximately 75 per cent of the plants will have red flowers and the rest white ones. This argument, which may be used for explanatory or for predictive purposes, is a statistical systematization; what it explains or predicts are the approximate percentages of red- and white-flowered plants in the sample; the "premises" by reference to which the specified percentages are shown to be highly probable include (1) the pertinent laws of genetics, some of which are statistical generalizations, whereas others are of strictly universal form; and (2) particular information of the kind mentioned above about the genetic make-up of the parent generation of the plants from which the sample is taken. (The genetic principles of strictly universal form include the laws that the colors in question are tied to specific genes; that the red gene is dominant over the white one; and various other general laws concerning the transmission, by genes, of the colors in question—or, perhaps, of a broader set of gene-linked traits. Among the statistical generalizations invoked is the hypothesis that the four possible combinations of color-determining genes—WW, WR, RW, RR—are statistically equiprobable in their occurrence in the offspring of two plants of the hybrid generation.) These premises may fairly be regarded as exhausting that part of the total available evidence that is relevant to the hypothesis about the composition of the sample. Similar considerations apply to the kind of argument that serves retrodictively to establish the time of manufacture of a wooden implement found at an archeological site when the estimate is based on the amount of radioactive carbon the implement contains. Again, in addition to statements of particular fact, the argument invokes hypotheses of strictly universal form as well as a statement, crucial to the argument at hand, concerning the rate of decay of radioactive carbon; this statement has the form of a statistical probability hypothesis.

Let us now examine one further example somewhat more closely. The statistical law that the half life of radon is 3.82 days may be invoked for a statistical explanation of the fact that within 7.64 days, a particular sample consisting of 10 milligrams of radon was reduced, by radioactive decay, to a residual amount falling somewhere within the interval from 2 to 3 milligrams; it could similarly be used for predicting a particular outcome of this kind. The gist of the explanatory and predictive argument is, briefly, this: The statement giving the half life of radon conveys two statistical laws, (i) that the statistical probability for an atom of radon to undergo radioactive decay within a period of 3.82 days is $\frac{1}{2}$, and (ii) that the decaying of different radon atoms constitutes statistically independent events. One further premise needed is the statement that the number of atoms in 10 milligrams of radon is enormously large (in excess of 10^{19}). As mathematical probability theory shows, the two laws in conjunction with this latter statement imply deductively that the statistical probability is exceedingly high that the mass of the radon atoms surviving after 7.64 days will not deviate from 2.5 milligrams by more than .5, i.e., that it will fall within the specified interval. More explicitly, the consequence deducible from the two statistical laws in conjunction with the information on the large number of atoms involved is another statistical law to this effect: The statistical probability is very high that the random experiment F of letting 10 milligrams of radon decay for 7.68 days will yield an outcome of kind G, namely a residual amount of radon whose mass falls within the interval from 2 to 3 milligrams. Indeed, the probability is so high that, according to the interpretation (9.2b), if the experiment F is performed just once, it is "practically certain" that the outcome will be of kind G. In this sense, it is rational on the basis of the given information to expect the outcome G to occur as the result of a single performance of F; and also in this sense, the information concerning the half life of radon and the large number of atoms involved in an experiment of kind F affords a statistical explanation or prediction of the occurrence of G in a particular performance of the experiment.[44]

[44] By reference to a physical theory that makes essential use of statistical systematization, Hanson [23], has recently advanced an interesting argument against the view

In the statistical systematization here outlined, the requirement of total evidence is satisfied at least in the broad sense that according to the total body of present scientific knowledge, the rate of radioactive decay of an element is independent of all other factors, such as temperature and pressure, ordinary magnetic and electric influences, and chemical interactions; so none of these need be taken into consideration in appraising the probability of the specified outcome.

Other statistical explanations offered in science for particular phenomena follow the same general pattern: To account for the occurrence of a certain kind of event under specified (e.g., experimental) conditions, certain laws or theories of statistical form are adduced, and it is shown that as a consequence of these, the statistical probability for an outcome of the specified kind under circumstances of the specified kind is extremely high, so that that outcome may be expected with practical certainty in any one case where the specified conditions occur. (For example, the probabilistic

that any explanation constitutes a potential prediction. According to Hanson, that view fits the character of the explanations and predictions made possible by the laws of Newtonian classical mechanics, which are deterministic in character; but it is entirely inappropriate for quantum theory, which is fundamentally nondeterministic. More specifically, Hanson holds that the laws of quantum theory do not permit the *prediction* of any individual quantum phenomenon P, such as the emission of a beta particle from a radioactive substance, but that "P can be completely *explained* ex post facto; one can understand fully just what kind of event occurred, in terms of the well-established laws of . . . quantum theory . . . These laws give the *meaning* of 'explaining single microevents'." (Hanson [23], p. 354; the italics are the quoted author's.) I quite agree that by reason of their statistical character, the laws of quantum theory permit the prediction of events such as the emission of beta particles by a radioactive substance only statistically and not with deductive-nomological certainty for an individual case. But for the same reason it is quite puzzling in what sense those laws could be held to permit a complete explanation ex post facto of the single event P. For if the explanans contains the statement that P has occurred, then the explanation is unilluminatingly circular; it might be said, at best, to provide a description of what in fact took place, but surely not an understanding of why it did; and to answer the question 'why?' is an essential task of explanation in the characteristic sense with which we have been, and will be, concerned throughout this essay. If, on the other hand, the explanans does not contain the statement that P has occurred, but only statements referring to antecedent facts plus the laws of quantum theory, then the information thus provided can at best show that an event of the kind illustrated by P—namely, emission of a beta particle—was highly probable under the circumstances; this might then be construed, in the sense outlined in the text, as constituting a probabilistic explanation for the occurrence of the particular event P. Thus, it still seems correct to say that an explanation in terms of statistical laws is also a potential prediction, and that both the explanation and the prediction are statistical-probabilistic in character, and provide no complete accounts of individual events in the manner in which deductive-nomological systematization permits a complete account of individual occurrences.

explanation provided by wave mechanics for the diffraction of an electron beam by a narrow slit is essentially of this type.)

Let us examine the logic of the argument by reference to a simple model case: Suppose that a statistical explanation is to be given of the fact that a specified particular sequence S of 10 successive flippings of a given coin yielded heads at least once—let this fact be expressed by the sentence h; and suppose furthermore we are given the statements that the statistical probabilities of heads and of tails for a flipping of the given coin both equal ½, and that the results of different flippings are statistically independent of each other. These statements might then be invoked to achieve the desired explanation; for jointly they imply that the probability for a set of 10 successive flippings of the given coin to yield heads at least once is $1 - (½)^{10}$, which is greater than .999. But this probability is still statistical in character; it applies to a certain kind of event (heads at least once) relative to a certain other kind of event (10 flippings of the given coin), but not to any individual event, such as the appearance of heads at least once in the particular unique set S of 10 flippings. If the statistical probability statement is to be used in explaining this latter event, then an additional principle is needed which makes statistical probabilities relevant to rational expectations concerning the occurrence of particular events.

One such principle is provided by the interpretation of a very high statistical probability as making it practically certain that the kind of outcome in question will occur in any one particular case (see (9.2b) above). This idea can be expressed in the following rule:

(11.1) On the information that the statistical probability $p(G, F)$ exceeds $1 - \epsilon$ (where ϵ is some very small positive number) and that b is a particular instance of F, it is practically certain that b is an instance of F.

Another way of giving statistical probability statements relevance for rational expectations concerning individual events would be to develop a system of inductive logic for languages in which statistical probability statements can be expressed. Such a system would assign, to any "hypothesis" h expressible in the language, a logical probability $c(h, e)$ with respect to any logically consistent evidence sentence e in that language. Choosing as evidence the sentence e_1, 'The statistical probability of obtaining heads at least once in a set of 10 flippings of this coin is $1 - (½)^{10}$, and S is a

particular set of such flippings,' and as hypothesis the sentence h_1, 'S yields heads at least once,' we would then obtain the logical probability conferred by e_1 on h_1. Now the systems of inductive logic presently available—by far the most advanced of which is Carnap's—do not cover languages rich enough to permit the formulation of statistical probability statements.[45] However, for the simple kind of argument under consideration here, it is clear that the value of the logical probability should equal that of the corresponding statistical probability, i.e., that we should have $c(h_1, e_1) = 1 - (½)^{10}$. Somewhat more generally, the idea may be expressed in the following rule:

(11.2) If e is the statement '$(p(G, F) = r) \cdot Fb$' and h is 'Gb,' then $c(h, e) = r$.

This rule is in keeping with the conception, set forth by Carnap, of logical probability as a fair betting quotient for a bet on h on the basis of e; and it accords equally with Carnap's view that the logical probability on evidence e of the hypothesis that a particular case b will have a specified property M may be regarded as an estimate, based on e, of the relative frequency of M in any class K of cases on which the evidence e does not report. Indeed, Carnap adds that the logical probability of 'Mb' on e may in certain cases be considered as an estimate of the statistical probability of M.[46] If, therefore, e actually contains the information that the statistical probability of M is r, then it seems clear that the estimate, on e, of that statistical probability, and thus the logical probability of 'Mb' on e, should be r as well.

The rules (11.1) and (11.2) may be regarded as schematizing at least simple kinds of statistical systematization. But, arguments conforming to those rules will constitute acceptable explanations or predictions only if they satisfy the principle of total evidence. For example, suppose that the total evidence e contains the information e_1 that F_1b and $p(G, F_1) = .9999$; then e_1 makes it practically certain that Gb; and yet it would not be acceptable as the premise of a statistical explanation or prediction of 'Gb' if e also contained the information, e_2, and F_2b and $p(G, F_2) = .0001$. By itself, e_2 makes it practically certain that b is not G; and if e consists of just e_1 and e_2, then the simple rule (11.2) does not enable us

[45] I learned from Professor Carnap, however, that in as yet unpublished work, his system of inductive logic has been extended to cover also statistical probability statements.

[46] See Carnap [7], pp. 168–175.

to assign a logical probability to 'Gb.' But suppose that, besides e_1 and e_2, e also contains e_3: '$p_3(G, F_1 \cdot F_2) = .9997$, and nothing else (i.e., nothing that is not logically implied by e_1, e_2, and e_3 in conjunction). Then it seems reasonable to say that the probability of 'Gb' on e should be equal, or at least close, to .9997. Similarly, if e contains just the further information that F_3b and $p(G, F_1 \cdot F_2 \cdot F_3) = .00002$ then the probability of 'Gb' on e should be close to .00002, and so on.

This consideration suggests the possibility of meeting the desideratum expressed at the end of Section 10 by the following rough substitute for the requirement of total evidence:

(11.3) *Rough criterion of evidential adequacy for simple statistical systematizations*: A statistical systematization of the simple type indicated in rules (11.1) and (11.2) may be regarded as satisfying the requirement of total evidence if it is based on the statistical probability of G within the narrowest class, if there is one, for which the total evidence e available provides the requisite statistical probability.[47] More explicitly, a statistical systematization with the premises 'Fb' and '$p(G, F) = r$' may be regarded as roughly satisfying the requirement of total evidence if the following conditions are met: (i) the total evidence e contains (i.e., explicitly states or deductively implies) those two premises; (ii) e implies[48] that F is a subclass of any class F* for which e contains the statement that F*b and in addition a statistical law (which must not be simply a theorem of formal probability theory)[49] stating the value of the probability $p(G, F^*)$. The classes F, F*, etc., are of course understood here simply as the classes of those elements which have the characteristics F, F*, etc.

[47] This idea is closely related to one used by Reichenbach (see [41], Sec. 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 and successful completion of a particular scheduled flight of a given commercial plane, be construed as the statistical probability which the *kind* of event considered (safe and successful 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 (this might be, for example, 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). Our working rule, however, assigns a probability to (a statement describing) a single event only if the total evidence specifies the value of the pertinent statistical probability; whereas Reichenbach's interpretation refers to the case where the total evidence provides a statistical report on a finite sample from the specified reference class (in our illustration, a report on the frequencies of safe completion in the finite class of similar flights undertaken so far); note that such a sample report is by no means equivalent to a statistical probability statement, though it may well suggest such a statement and may serve as supporting evidence for it. (On this point, cf. also Sec. 7 of the present essay.)

[48] This requirement of implication serves to express the idea that F is the narrowest class of which b is *known* (namely, as a consequence of the total evidence) to be an element.

Condition (ii) might be liberalized by the following qualification: F need not be the narrowest class of the kind just specified; it suffices if e implies that within any subclass of F to which e assigns b, the statistical probability of G is the same as in F. For example, in the prediction, considered above, of the residual mass of radon, the total information available may well include data on temperature, pressure, and other characteristics of the given sample s: In this case, e assigns the particular event under study to a considerably narrower class than the class F of cases where a 10 milligram sample of radon is allowed to decay for 7.64 days. But the theory of radioactivity, likewise included in e, implies that those other characteristics do not affect the probability invoked in the prediction; in other words, the statistical probability of decay in the corresponding subclasses of F is the same as in F itself.

The working rule suggested here would also avoid an embarrassment which the general requirement of total evidence creates for the explanatory use of statistical systematizations. Suppose, for example, that an individual case b has been found to have the characteristic G (or to belong to the class G); and consider a proposed explanation of Gb by reference to the statements 'Fb' and '$p(G, F) = .9999$.' Even assuming that nothing else is known, the total evidence then includes, in addition to these latter two statements, the sentence 'Gb.' Hence if we were strictly to enforce the requirement of total evidence, then the explanans, by virtue of containing the explanandum, would trivially imply the later without benefit of any statistical law, and would confer upon it the logical probability 1. Thus, no nontrivial inductive explanation would be possible for any facts or events that are known (reported by e) to have occurred. This consequence cannot be avoided by the convention that e with the explanandum statement omitted is to count as total evidence for the statistical explanation of an event known to have occurred; for despite its apparent clarity, the notion of omitting the explanandum statement from e does not admit of a precise logical explication. It is surely not a matter of just deleting the

[49] Statistical probability statements which are theorems of mathematical probability theory cannot properly be regarded as affording an explanation of empirical subject matter. The condition will prove significant in a context to be discussed a little later in this section.

explanandum sentence from e, for the total evidence can always be so formulated as not to contain that sentence explicitly; for example, 'Gb' may be replaced by the two sentences 'Gb ∨ Fb' and 'Gb ∨ — Fb.'

On the other hand, the working rule would circumvent the difficulty. For even though, in the illustration, e contains 'Gb,' the rule qualifies the statistical explanation of 'Gb' by means of 'Fb' and 'p(G, F) = .9999' alone as satisfying the requirement of total evidence. For the statistical law invoked here specifies the probability of G for the narrowest reference class to which e assigns b, namely the class F. (To be sure, e also assigns b to the narrower reference class F · G, for which clearly p(G, F · G) = 1. It will be reasonable to say that e (trivially) contains this latter statement since it is simply a logical consequence of the measure-theoretical postulates for statistical probability. But precisely for this reason, the statement 'p(G, F · G) = 1' is not an empirical law; hence, under the working rule, this part of the content of e need not be taken into consideration.)[50]

But while a rule such as (11.3) does seem in accord with the rationale of scientific arguments intended to explain or to predict individual occurrences by means of statistical laws, it offers no more than a rough working principle, which must be used with caution and discretion. Suppose, for example, that the total evidence e consists of the statements 'Fb,' 'Hb,' 'p(G, F) = .9999,' and a report on 10,000 individual cases other than b, to the effect that all of them were H and non-G. Then the statistical argument with 'Fb' and 'p(G, F) = .9999' as its premises and 'Gb' as its con-

[50] While here our rule permits us to disregard, as it were, the occurrence of the explanandum in the total evidence, this is not so in all cases. Suppose, for example, that the total evidence e consists of the following sentences: e_1: 'p(G, F) = .4'; e_2: 'p(G,(G ∨ H) · F) = .9999'; e_3: 'Fb'; e_4: 'Gb.' Here again, e assigns b to the class F · G, for which p(G, F · G) = 1; as before, we may disregard this narrowest reference class. But e implies as well that b belongs to the class (G ∨ H) · F, which is the narrowest reference class relative to which e also specifies an empirical probability for G. Hence under our rule the statistical systematization with the premises e_2 and '(Gb ∨ Hb) · Fb' and with the conclusion 'Gb' satisfies the requirement of total evidence (whereas the argument with e_1 and e_3 as premises and 'Gb' as conclusion does *not*). Thus, we have here an argument that statistically explains b's being G by reference to b's being (G ∨ H) · F, though to establish this latter fact, we made use of the sentence 'Gb.' In this case, then, our rule does not allow us to disregard the occurrence of the explanandum in the total evidence.

The logical situation illustrated here seems to be analogous to that described by Scriven [50], in reference to causal explanation. Scriven points out that when we causally explain a certain event by reference to certain antecedent circumstances, it may happen that practically the only ground we have for assuming that those explanatory antecedents were in fact present is the information that the explanandum event did occur. Similarly, in our illustration, the information that b is G provides the ground for the assertion that b has the explanatory characteristic (G ∨ H) · F.

clusion would qualify, under the rule, as meeting the requirement of total evidence; but even though e does not state the statistical probability of G relative to H, its sample statistics on 10,000 cases of H, in conjunction with the statement that b is H, must surely cast serious doubt upon the acceptability of the proposed statistical argument as an explanation or prediction of Gb. Hence, the information relevant to 'Gb' that is provided by e cannot generally and strictly be identified with the information provided by e concerning the statistical probability of G in the narrowest available reference class; nor, of course, can the logical probability of 'Gb' on e be strictly equated with the statistical probability of G in that narrowest reference class. Thus, as a general condition for a statistical systematization that is to be not only a formally correct argument (a potential systematization) but a scientifically acceptable one, the requirement of total evidence remains indispensable.

12. *On Criteria of Rational Credibility.*

Besides the requirement of total evidence, there is a further condition which it might seem any statistical systematization ought to satisfy if it is to qualify as an adequate explanation, prediction, or retrodiction; namely, that the information contained in its "premises" e should provide so strong a presumption in favor of the "conclusion" h as to make it rational, for someone whose total evidence is e, to believe h to be true, or, as I will also say, to include h in the set of statements accepted by him as presumably true. In a deductive-nomological systematization, the premises afford such presumption in an extreme form: they logically imply the conclusion; hence someone whose system of accepted statements includes those premises has the strongest possible reason to accept the conclusion as well.

Thus, the study of inductive generalization gives rise to the question whether it is possible to formulate criteria for the rational acceptability of hypotheses on the basis of information that provides strong, but not conclusive, evidence for them.

I will first construe this question in a quite general fashion without limiting it specifically to the case where the supporting information provides the premises of a statistical systematization. Toward the end of this section I will return to this latter, special case.

Let us assume that the total body of scientific knowledge at a given time t can be represented by a set K_t, or briefly K, whose elements are all the statements accepted as presumably true by the scientists at time t. The class K will contain statements describing particular events as well as assertions of statistical and universal law and in addition various theoretical statements. The membership of K will change in the course of time; for as a result of continuing research, additional statements come to be established, and thus accepted into K; while others, formerly included in K, may come to be disconfirmed and then eliminated from the system.

We can distinguish two major ways in which a statement may be accepted into K: *direct acceptance*, on the basis of suitable experiences or observations, and *inferential acceptance*, by reference to previously accepted statements. An observer who records the color of a bird or notes the reading of an instrument accepts the corresponding statements directly, as reporting what he immediately observes, rather than as hypotheses whose acceptability is warranted by the fact that they can be inferred from other statements, which have been antecedently accepted and thus are already contained in K. Inferential acceptance may be either deductive or (strictly) inductive, depending on whether the statement in question is logically implied or only more or less highly supported by the previously accepted statements.

This schematic model does not require, then, that the statements representing scientific knowledge at a given time be true; rather, it construes scientific knowledge as the totality of beliefs that are accepted at a given time as warranted by appropriate scientific procedures. I will refer to this schematization as the *accepted-information model of scientific knowledge*.

Now, we have to consider the rules of acceptance or rejection which regulate membership in K. In its full generality, this question calls for a comprehensive set of principles for the formulation, test, and validation of scientific hypotheses and theories. In the context of our investigation, however, it will suffice to concentrate on some general rules for indirect acceptance; the question of criteria for direct acceptance, which would bear on standards for observational and experimental procedures, is not relevant to the central topic of this essay.

The rules to be discussed here may be considered as stating certain necessary conditions of rationality in the formation of beliefs. One very obvious condition of this kind is the following:

(CR1) Any logical consequence of a set of accepted statements is likewise an accepted statement; or, K contains all logical consequences of any of its subclasses.

The reason for this requirement is clear: If an investigator believes a certain set of statements, and thus accepts them as presumably true, then, to be rational, he has to accept also their logical consequences because any logical consequence of a set of true statements is true.

Note that (CR1) does not express a rule or principle of logic but rather a maxim for the rational *application* of the rules of deductive logic. These rules, such as *modus ponens* or the rules of the syllogism, simply indicate that if sentences of a specified kind are true, then so is a certain other sentence; but they say nothing at all about what it is rational to believe. Another rule is the following:

(CR2) The set K of accepted statements is logically consistent.

Otherwise, by reason of (CR1), K would also contain, for every one of its statements, its contradictory. This would defeat the objective of science of arriving at a set of presumably true beliefs (if a statement is presumably true, its contradictory is not); and K could provide no guidance for expectations about empirical phenomena since whatever K asserted to be the case it would also assert not to be the case.

(CR3) The inferential acceptance of any statement h into K is decided on by reference to the total system K (or by reference to a subset K′ of it whose complement is irrelevant to h relative to K′).

This is simply a restatement of the requirement of total evidence. As noted earlier, it is automatically satisfied in the case of deductive acceptance.

Now we must look for more specific rules of inferential acceptance. The case of deductive acceptance is completely settled by (CR1), which makes it obligatory for rational procedure to accept all statements that are deductively implied by those already accepted. Can analogous rules be specified for rational inductive acceptance? Recent developments in the theory of inductive procedures suggest that this question might best be considered as a special case of the general problem of establishing criteria of rationality for choices between several alternatives; in the case at hand, the choice would be that of accepting a proposed new statement h into K, rejecting it (in the strong sense of accepting its contradictory), or leaving the decision in suspense (i.e., accepting neither h nor its contradictory).

I will consider the problem first on the assumption that a system of inductive logic is available which, for any hypothesis h and for any logically consistent "evidence" sentence e, determines the logical probability, or the degree of confirmation, c(h, e), of h relative to e.

The problem of specifying rational rules of decision may now be construed in the following schematic fashion: An agent X has to choose one from among n courses of action, $A_1, A_2, \ldots, A_n$, which, on the total evidence e available to him, are mutually exclusive and jointly exhaust all the possibilities open to him. Each of these may eventuate, with certain probabilities (some of which may be zero), in any one of m outcomes, $O_1, O_2, \ldots, O_m$, which, on the evidence e, are mutually exclusive and exhaustive. The agent's decision to choose a particular course of action, say A_k, will be rational only if it is based on a comparison of its probable consequences with those of the alternative choices that are open to him. For such a comparison, inductive logic would provide one important tool. Let $a_1, a_2, \ldots, a_n$ be statements to the effect that X follows course of action $A_1, A_2, \ldots, A_n$, respectively; and let $o_1, o_2, \ldots, o_m$ be statements asserting the occurrence of $O_1, O_2, \ldots, O_m$, respectively. Then the probability, relative to e, for a proposed course of action, say A_j, to yield a specified outcome, say O_k, is given by $c(o_k, e \cdot a_j)$. The principles of the given system of inductive logic would determine all these probabilities, but they would not be sufficient to determine a rational course of action for X. Indeed, rationality is a relative concept; a certain decision or procedure can be qualified as rational only relative to some objective, namely by showing, generally speaking, that the given decision or procedure offers the optimal prospect of attaining the stated objective.

One theoretically attractive way of specifying such objectives is to assume that for X each of the outcomes $O_1, O_2, \ldots, O_n$ has a definite value or disvalue, which is capable of being represented in quantitative terms by a function assigning to any given outcome, say O_k, a real number u_k, the utility of O_k for X at the time in question. The idea of such a utility function raises a variety of problems which cannot be dealt with here, but which have been the object of intensive discussion and of much theoretical as well as experimental research.[51] The utility function, together with the probabilities just mentioned, determines the expectation value, or the probability estimate, based on e, of the utility attached to A_j for X:

$$(12.1) \quad u'(A_j, e) = c(o_1, e \cdot a_j) \cdot u_1 + \ldots + c(o_m, e \cdot a_j) \cdot u_m.$$

[51] For details and further bibliographic references, see, for example, Neumann and Morgenstern [36]; Savage [45]; Luce and Raiffa [33]; Carnap [7], par. 51.

In the context of our schematization, the conception of rationality of decision as relative to some objective can now be taken into account in a more precise form; this is done, for example, in the following rule for rational choice, which was proposed by Carnap:

Rule of maximizing the estimated utility: In the specified circumstances, X acts rationally if he chooses a course of action, A_j, for which the expectation value of the utility is maximized, i.e., is not exceeded by that associated with any of the alternative courses of action.[52]

I will now attempt to apply these considerations to the problem of establishing criteria of rational inductive acceptance. The decision to accept, or to reject, a given hypothesis, or to leave it in suspense might be considered as a special kind of choice required of the scientific investigator. This conception invites an attempt to obtain criteria of rational inductive belief by applying the rule of maximizing the expected utility to this purely scientific kind of choice with its three possible "outcomes": K enlarged by the contemplated hypothesis h; K enlarged by the contradictory of h; K unchanged. But what could determine the utilities of such outcomes?

The pursuit of knowledge as exemplified by pure scientific inquiry, by "basic research" not directly aimed at any practical applications with corresponding utilities, is often said to be concerned with the discovery of truth. This suggests that the acceptance of a hypothesis might be considered a choice as a result of which either a truth or a falsehood is added to the previously established system of knowledge. The problem then is to find a measure of the purely scientific utility, or, as I will say, the *epistemic utility*, of such an addition.

It seems reasonable to say that the epistemic utility of adding h to K depends not only on whether h is true or false but also on how much of what h asserts is new, i.e., goes beyond the information already contained in K. Let k be a sentence which is logically implied by K, and which in turn implies every sentence in K, just as the conjunction of the postulates in a finite axiomatization of geometry implies all the postulates and theorems of geometry. Then k has the same informational content as K. Now,

[52] Cf. Carnap [7], p. 269; the formulation given there is "Among the possible actions choose that one for which the estimate of the resulting utility is a maximum." Carnap proposes this rule after a critical examination, by reference to instructive illustrations, of several other rules for rational decision that might seem plausible (*ibid.*, Secs. 50, 51).

the common content of two statements is expressed by their disjunction, which is the strongest statement logically implied by each of them. Hence, the common content of h and K is given by $h \vee k$. But h is equivalent to $(h \vee k) \cdot (h \vee -k)$, where the two component sentences in parentheses have no common content: their disjunction is a logical truth. Hence that part of the content of h which goes beyond the information contained in K is expressed by $(h \vee -k)$. To indicate *how much* is being asserted by this statement, we make use of the concept of a content measure for sentences in a (formalized) language L. By a *content measure function* for a language L we will understand a function m which assigns, to every sentence *s* of L, a number m(s) in such a way that (i) $0 \leqq m(s) \leqq 1$; (ii) $m(s) = o$ if and only if s is a logical truth of L; (iii) if the contents of s_1 and s_2 are mutually exclusive—i.e., if the disjunction $s_1 \vee s_2$ is a logical truth of L—then $m(s_1 \cdot s_2) = m(s_1) + m(s_2)$.[53] (If these requirements are met, then m can readily be seen to satisfy also the following conditions: (iv) $m(s) = 1 - m(-s)$; (v) if s_1 logically implies s_2, then $m(s_1) \geqq m(s_2)$; (vi) logically equivalent sentences have equal measures.)

Let m be a content measure function for an appropriately formalized language suited to the purposes of empirical science. Then, in accordance with the idea suggested above, it might seem plausible to accept the following:

(12.2) *Tentative measure of epistemic utility*: The epistemic utility of accepting a hypothesis h into the set K of previously accepted scientific statements is $m(h \vee -k)$ if h is true, and $-m(h \vee -k)$ if h is false; the utility of leaving h in suspense, and thus leaving K unchanged, is 0.

The rule of maximizing the estimated utility now qualifies the decision to accept a proposed hypothesis as epistemically rational if the probability estimate of the corresponding utility is at least as great as the estimates attached to the alternative choices. The three estimates can readily be computed. The probability, on the basis of K, that the proposed hypothesis h is true is c(h, k), and that it is false, $1 - c(h, k)$. Denoting the three alternative actions of accepting h, rejecting h, and leaving h in suspense

[53] Content measures satisfying the specified conditions can readily be constructed for various kinds of formalized languages. For a specific measure function applicable to any first-order functional calculus with a finite number of predicates of any degrees, and a finite universe of discourse, see *SLE*, par. 9, or Carnap and Bar-Hillel [9], Sec. 6.

by 'A,' 'R,' 'S,' respectively, we obtain the following formulas for the estimated utilities attached to these three courses of action:

$$(12.3a)\ u'(A, k) = c(h, k) \cdot m(h \vee -k) - (1 - c(h, k)) \cdot m(h \vee -k) = m(h \vee -k) \cdot (2c(h, k) - 1).$$

Analogously, considering that rejecting h is tantamount to accepting —h, which goes beyond K by the assertion $-h \vee -k$, we find

$$(12.3b)\ u'(R, k) = m(-h \vee -k) \cdot (1 - 2c(h, k)).$$

Finally, we have

$$(12.3c)\ u'(S, k) = 0.$$

Now the following can be readily verified:[54]

(i) If $c(h, k) = 1/2$, then all three estimates are zero;
(ii) If $c(h, k) > 1/2$, then u'(A, k) exceeds the other two estimates;
(iii) If $c(h, k) < 1/2$, then u'(R, k) exceeds the other two estimates.

Hence, the principle of maximizing the estimated utility leads to the following rule:

(12.4) *Tentative rule for inductive acceptance*: Accept or reject h, given K, according as $c(h, k) > 1/2$ or $c(h, k) < 1/2$; when $c(h, k) = 1/2$, h may be accepted, rejected, or left in suspense.

It is of interest to note that the principle of maximizing the estimated utility, in conjunction with the measure of epistemic utility specified in (12.2), implies this rule of acceptance quite independently of whatever particular inductive probability function c and whatever particular measure function m might be adopted.

Unfortunately, the criteria specified by this rule are far too liberal to be acceptable as general standards governing the acceptance of hypotheses in pure science. But this does not necessarily mean that the kind of approach attempted here is basically inadequate: the fault may well lie with the oversimplified construal of epistemic utility. It would therefore seem a problem definitely worth further investigation whether a modified version of the concept of epistemic utility cannot be construed which, via

[54] We have

$$u'(A, k) - u'(R, k) = (2c(h, k) - 1) \cdot (m(h \vee -k) + m(-h \vee -k)).$$

Since m is nonnegative, the second factor on the right could be 0 only if both of its terms were 0. But this would require $h \vee -k$ as well as $-h \vee -k$ to be logically true, in which case k would logically imply both h and —h; and this is precluded by the consistency requirement, (CR2), for K. Hence, u'(A, k) exceeds u'(R, k) or is exceeded by it according as c(h, k) is greater or less than 1/2; and whichever of the two estimates is the greater will also be positive and thus greater than u'(S, k).

the principle of maximizing estimated utility, will yield a more satisfactory rule for the inductive acceptance or rejection of hypotheses in pure science. Such an improved measure of epistemic utility might plausibly be expected to depend, not only on the change of informational content, but also on other changes in the total system of accepted statements which the inductive acceptance of a proposed hypothesis h would bring about. These would presumably include the change in the simplicity of the total system, or, what may be a closely related characteristic, the change in the extent to which the theoretical statements of the system would account for, or systematize, the other statements in the system, in particular those which have been directly accepted as reports of previous observational or experimental findings. As yet, no fully satisfactory general explications of these concepts are available, although certain partial results have been obtained.[55] And even assuming that the concepts of simplicity and degree of systematization can be made explicit and precise, it is yet another question whether the notion of epistemic utility permits a satisfactory explication, which can serve as a basis for the construction of rules of inductive acceptance.

We will now consider briefly an alternative construal of scientific knowledge, which would avoid the difficulties just outlined: it will be called the *pragmatist* or *instrumentalist model*. Let us note, first of all, that the epistemic utilities associated with the decision inductively to accept (or to reject, or to leave in suspense) a certain hypothesis would have to represent "gains" or "losses" as judged by reference to the objectives of "pure" or "basic" scientific research; in contradistinction to what will be called here *pragmatic utilities*, which would represent gains or losses in income, prestige, intellectual or moral satisfaction, security, and so forth, that may accrue to an individual or to a group as a result of "accepting" a proposed hypothesis in the practical sense of basing some course of action on it. Theories of rational decision making have usually been illustrated by, and applied to, problems in which the utilities are of this pragmatic kind, as for example, in the context of quality control. The hypotheses that have to be considered in that case concern the items produced by a certain technological process during a specified time; e.g., vitamin capsules which

[55] For an illuminating discussion of the concept of simplicity of a total system of statements, see Barker [2] (especially Chs. 5 and 9); also see the critical survey by Goodman [22]. One definition (applicable only to formalized languages of rather simple structure) of the systematizing power of a given theory with respect to a given class of data has been proposed in *SLE*, Secs. 8 and 9.

must meet certain standards, or tablets containing a closely specified amount of a certain toxic ingredient, or ball bearings for whose diameter a certain maximum tolerance has been fixed, or light bulbs which must meet various specifications. The hypothesis under test will assert, in the simplest case, that the members of the population (e.g., the output produced by a given industrial plant in a week) meet certain specified standards (e.g., that certain of their quantitative characteristics fall within specified numerical intervals). The hypothesis is tested by selecting a random sample from the total population and examining its members in the relevant respects. The problem then arises of formulating a general decision rule which will indicate, for every possible outcome of the test, whether on the evidence afforded by that outcome the hypothesis is to be accepted or rejected. But what is here referred to as acceptance or rejection of a hypothesis clearly amounts to adopting or rejecting a certain practical course of action (e.g., to ship the ball bearings to the distributors, or to reprocess them). In this kind of situation, we may distinguish four possible "outcomes": the hypothesis may be accepted and in fact true, rejected though actually true, accepted though actually false, or rejected and in fact false. To each of these outcomes there will be attached a certain positive or negative utility, which in cases of the kind considered might be represented, at least approximately, in monetary terms. Once such utilities have been specified, it is possible to formulate decision rules which will indicate for every possible outcome of the proposed testing procedure whether, on the evidence provided by the outcome and in consideration of the utilities involved, the hypothesis is to be accepted or to be rejected. For example, the principle of maximizing estimated utilities affords such a rule, which presupposes, however, that a suitable inductive logic is available which assigns to any proposed hypothesis h, relative to any consistent "evidence" statement e, a definite logical probability, c(h, e).

Alternatively, there have been developed, in mathematical statistics and in the theory of games, certain methods of arriving at decision rules which do not require any such general concept of inductive or logical probability. These methods are limited to certain special types of hypotheses and evidence sentences; normally, their application is to hypotheses in the form of probability statements (statistical generalizations), and to evidence sentences in the form of reports on statistical findings in finite samples. One of the best known methods of this kind is based on the minimax principle. This method uses the concept of probability only in its statistical form. It is intended to select the most rational from among various possible rules that might be followed in deciding on the acceptance or rejection of a proposed hypothesis h in consideration of (i) the results of a specified kind of test and (ii) the utilities assigned to the possible "outcomes" of accepting or rejecting the hypothesis. Briefly, the minimax principle directs that we adopt, from among the various possible decision rules, one that minimizes the maximum risk, i.e., one for which the largest of the (statistically defined) probability estimates of the losses that might be incurred in the given context as a result of following this rule is no greater than the largest of the corresponding risks (loss estimates) attached to any of the alternative decision rules.[56]

Clearly, the minimax principle is not itself a decision rule, but rather a metarule specifying a standard of adequacy, or of rationality, for decision rules pertaining to a suitably characterized set of alternative hypotheses, plus testing procedure, and a given set of utilities.[57]

But whatever decision rules, or whatever general standards for the choice of decision rules, may be adopted in situations of the kind referred to here, the crucial point remains that the pragmatic utilities involved, and thus the decision dictated by the rule once the test results are given, will depend on, and normally vary with, the kind of action that is to be based upon the hypothesis. Consider, for example, the hypothesis that all of the vials of vaccine produced during a given period of time by a pharmaceutical firm meet certain standards of purity; and suppose that a test has been performed by analyzing the vials in a random sample. Then the gains or losses to be expected from correct or incorrect assumptions as to the truth of the hypothesis will depend on the action that is intended, for example, on whether the vaccine is to be administered to humans or to animals. By reason of the different utilities involved, a given decision rule—be it the rule of maximizing estimated utility or a rule selected in accordance with the minimax, or a similar, standard—may then well specify, on one and the same evidence, that the hypothesis is to be rejected in the case of application to human subjects, but accepted if the application is to be to animals.

Clearly then, in cases of this kind we cannot properly speak of a decision to accept or reject a hypothesis per se; the decision is rather to adopt one of two (or more) alternative courses of action. Moreover, it is not even clear on what grounds the acceptance or rejection of this hypothesis per se, on the given evidence, could be justified—unless it is possible to specify a satisfactory concept of epistemic utility, whose role for the decisions of pure science would be analogous to that of pragmatic utility in decisions concerning actions based on scientific hypotheses.

Some writers on the problems of rational decision have therefore argued that one cannot strictly speak of a decision to accept a scientific hypothesis, and that the decisions in question have to be construed as concerning choices of certain courses of action.[58] A lucid presentation and defense of

[56] The minimax principle was proposed and theoretically developed by A. Wald; see especially his book [54]. A lucid and stimulating less technical account and appraisal of the minimax method, of special interest from a philosophical point of view, is given in Braithwaite [3], Ch. VII. Recent very clear presentations of the fundamentals of minimax theory, plus critical comments and further developments, may be found, for example, in Savage [45], and in Luce and Raiffa [33]. Carnap [7], par. 98, gives an instructive brief comparison of those methods of estimation which are based on inductive logic with those which, like the minimax method, have been developed within the framework of statistical probability theory, without reliance on a general inductive logic.

[57] The standard set up by the minimax principle is by no means the only possible standard of rationality that can be proposed for decision rules in problem situations of the kind referred to here; and indeed, the minimax standard has been criticized in certain respects, and alternatives to it have been suggested by recent investigators. For details, see, for example, Savage [45], Ch. 13; Luce and Raiffa [33], Ch. 13. In an article which includes a lucid examination of the basic ideas of the minimax principle, R. C. Jeffrey points out that in applying this principle the experimenter acts on the assumption that this is the worst of all possible worlds for him; thus "the minimax criterion is at the pessimistic end of a continuum of criteria. At the other end of this continuum is the 'minimin' criterion, which advises each experimenter to minimize his minimum risk. Here each experimenter acts as if this were the best of all possible worlds for him." (Jeffrey [28], p. 244.)

[58] See, for example, De Finetti [12], p. 219; Neyman [37], pp. 259–260. Savage [45], Ch. 9, Sec. 2, strongly advocates a "behavioralistic" as opposed to a "verbalistic" outlook on statistical decision problems; he argues that these problems are concerned with acts rather than with "assertions" (i.e., of scientific hypotheses). However, he grants the possibility of considering an "assertion" as a special kind of behavioral act and thus does not rule out explicitly the possibility of speaking of the acceptance—as presumably the same thing as "assertion"—of hypotheses in science. Savage here also makes some suggestive though all too brief remarks on the subtle practical consequences resulting from the assertion of a hypothesis in pure science (such as that the velocity of light is between 2.99×10^{10} and 3.01×10^{10} cm/sec); those consequences would presumably have to be taken into account, from his behavioralistic point of view, in appraising the utilities attached to the acceptance or rejection of purely scientific hypotheses. But Savage stresses that "many problems described according to the verbalistic outlook as calling for decisions between assertions really call only for decisions between much more down-to-earth acts, such as whether to issue single—or double—edged razors to an army . . ." (*loc. cit.*, p. 161). A distinction similar to that drawn by Savage is considered by Luce and Raiffa [33], who contrast "classical statistical inference" with "modern statistical decision theory" (Ch. 13, Sec. 10). In this context, the authors briefly consider the question of how to appraise the losses from falsely rejecting or accepting a scientific research hypothesis. They suggest that no such evaluation appears possible, but conclude with a remark that seems to

this point of view has been given by R. C. Jeffrey, who accordingly arrives at the conclusion that the scientist's proper role is to provide the rational agents of his society with probabilities for hypotheses which, on the more customary account, he would be described as simply accepting or rejecting.[59]

This view, then, implies a rejection of the accepted-information model of scientific knowledge and suggests an alternative which might be called a tool-for-optimal-action model, or, as I said earlier, an instrumentalist model of scientific knowledge. This label is meant to suggest the idea that whether a hypothesis is to be accepted or not will depend upon the sort of action to be based on it, and on the rewards and penalties attached to the possible outcomes of such action. An instrumentalist model might be formulated in different degrees of refinement. A very simple version would represent the state of scientific knowledge at a given time t by a set D, or more explicitly, D_t, of directly accepted statements, plus a theory of inductive support which assigns to each proposed hypothesis, or to at least some of them, a certain degree of support relative to D_t. Like K in the accepted-information model, D_t would be assumed to be logically consistent and to contain any statement logically implied by any of its own subsets. But no statement other than those in D_t, however strongly confirmed by D_t, would count as accepted, or as belonging to the scientific knowledge at the given time. Rather, acceptance would be understood pragmatically in the context of some contemplated action, and a decision would then depend on the utilities involved.

If, in particular, the theory of inductive support assumed here is an inductive logic in Carnap's sense, then it will assign a degree of confirmation c(h, e) to any statement h relative to any logically consistent statement e in the language of science, which we assume to be suitably chosen and formalized. In this case, scientific knowledge at a given time t might be represented by a functional k_t assigning to every sentence S that is expressible in the language of science a real-number value, $k_t(S)$, which lies between 0 and 1 inclusive. The value $k_t(S)$ would simply be the logical probability of S relative to D_t; in particular, for any S included in or logically implied by D_t, $k_t(S)$ would be 1; for any S logically incompatible with D_t, $k_t(S)$ would be 0. Temporal changes of scientific knowledge would be reflected by changes of k_t, and thus by changes in the numbers assigned to some of the sentences in the language of science.

hint at what I have called the concept of epistemic utility: ". . . if information is what is desired, then this requirement should be formalized and attempts should be made to introduce the appropriate information measures as a part of the loss structure. This hardly ends the controversy, however, for decision theorists are only too aware that such a program is easier suggested than executed!" (*Loc. cit.*, p. 324.)

[59] Jeffrey [28], p. 245.

But clearly, this version of a pragmatist model is inadequate: It construes scientific knowledge as consisting essentially of reports on what has been directly observed, for the formal theory of inductive probability which it presupposes for the appraisal of other statements would presumably be a branch of logic rather than of empirical science. This account of science disregards the central importance of theoretical concepts and principles for organizing empirical data into patterns that permit explanation and prediction. So important is this aspect of science that theoretical considerations will often strongly influence the decision as to whether a proposed report on some directly observed phenomenon is to be accepted: What is a fact is to some extent determined by theory. In this respect the notion of a system D_t of statements which are accepted directly and independently of theoretical considerations, and by reference to which the rational credibility of all other scientific statements is adjudged, is a decided oversimplification. And it is an oversimplification in yet another respect: In theoretically advanced disciplines, many of the terms that the experimenter would use to record his observations, and thus to formulate his directly accepted statements, belong to the theoretical vocabulary rather than to that of everyday observation and description; and the appropriate theoretical framework has to be presupposed if those statements are to make sense.

Another inadequacy of the model lies in the assumption that any individual hypothesis that may be proposed in the language of science can be assigned a reasonable degree of confirmation by checking it against the total set D_t of statements that have been directly accepted, for the test of any even moderately advanced scientific hypothesis will require the assumption of other hypotheses in addition to observational findings. As Duhem emphasized so strongly, what can be tested experimentally is never a single theoretical statement, but always a comprehensive and complexly interconnected body of statements.

If we were to try to construct a pragmatist model on the basis of statistical decision theory, the difficulties would become even greater; for this theory, as noted earlier, eschews the assignment of degrees of support to statements relative to other statements. Hence, here, the scientist would have to assume the role of a consultant who, in a limited class of experimental contexts, provides decisions concerning the acceptance or rejection of certain statistical hypotheses for the guidance of action, provided that the pertinent utilities have been furnished to him.

At present, I do not know of a satisfactory *general* way of resolving the issue between the two conceptions of science which are schematized by our two models. But the preceding discussion of these models does seem to suggest an answer to the question raised at the beginning of this section, namely, whether it should be required of a statistical explanation, prediction, etc. in science that its premises make its conclusion rationally acceptable.

The preceding considerations seem to indicate that it would be pointless to formulate criteria of acceptability by reference to pragmatic utilities; for we are concerned here with purely theoretical (in contrast to applied) explanatory and predictive statistical arguments. We might just add the remark that criteria of rational acceptability based on pragmatic utilities might direct us to accept a certain predictive hypothesis, even though it was exceedingly improbable on the available evidence, on the ground that, if it were true, the utility associated with its adoption would be exceedingly large. In other words, if a decision rule of this kind, which is based on statistical probabilities and on an assignment of utilities, singles out, on the basis of evidence e, a certain hypothesis h from among several alternatives, then what is qualified as rational is, properly speaking, not the decision to believe h to be true, but the decision to *act* in the given context *as if* one believed h to be true even though e may offer very little support for that belief.

The rational credibility of the conclusion, in a sense appropriate to the purely theoretical, rather than the applied, use of statistical systematizations will thus have to be thought of as represented by a suitable concept of inductive support (perhaps in conjunction with a concept of epistemic utility). And at least for the types of statistical systemization covered by rule (11.1) or (11.2), the statistical or logical probability specified in the argument itself may serve as an indicator of inductive support; the requirement of high credibility for the conclusion can then be met by requiring, in the case of (11.1), that ϵ be sufficiently small, and in the case of (11.2), that r be sufficiently large.

But the notions "sufficiently small" and "sufficiently large" invoked here cannot well be construed as implying the existence of some fixed

probability value, say r*, such that a statistical systematization will meet the requirement of rational credibility just in case the probability associated with it exceeds r*: The standards of rational credibility will vary with the context in which a statistical systematization is used.[60] It will therefore be more satisfactory, for an explication of the logic of statistical explanation, prediction, and similar arguments, explicitly to construe *statistical systematization as admitting of degrees*: The evidence e adduced in an argument of this kind may then be said to explain, or predict, or retrodict, or generally to systematize its "conclusion" h to degree r, where r is the inductive support that e gives to h. In this respect, statistical systematization differs fundamentally from its deductive-nomological counterpart: In a deductive-nomological systematization, the explanandum follows logically from the explanans and thus is *certain* relative to the latter; no higher degree of rational credibility (relative to the information provided by the premises) is possible, and anything less than it would vitiate the claim of a proposed argument to constitute a deductive systematization.

13. *The Nonconjunctiveness of Statistical Systematization.*

Another fundamental difference between deductive and statistical systematization is this: Whenever a given explanans e deductively explains each of n different explananda, say $h_1, h_2, \ldots, h_n$, then e also deductively explains their conjunction; but if an explanans e statistically explains each of n explananda, $h_1, h_2, \ldots, h_n$ to a positive degree, however high, it may still attribute a probability of zero to their conjunction. Thus, e may statistically explain (or analogously, predict, retrodict, etc.) very strongly whatever is asserted by each of n hypotheses, but not at all what is asserted by them conjointly: statistical systematization is, in this sense, nonadditive, or nonconjunctive (whereas deductive systematization is additive, or conjunctive). This point can be stated more precisely as follows:

(13.1) *Nonconjunctiveness of statistical systematization*: For any probability value p*, however close to 1, there exists a set of statisti-

[60] Even decision rules of the kind discussed earlier, which are formulated by reference to certain probabilities and utilities, provide only a comparative, not an absolute (classificatory) concept of rationality, i.e., they permit, basically, a *comparison* of any two in the proposed set of alternative choices and determine which of them, if any, is more rational than the other; thus, they make it possible to single out a most rational choice from among a set of available alternatives. But they do not yield a classificatory criterion which would characterize any *one* of the alternatives, either as rational or as nonrational in the given context.

cal systematizations which have the same "premise" e, but different "conclusions," $h_1, h_2, \ldots, h_n$, such that e confers a probability of at least p^* on every one of these conclusions but the probability zero on their conjunction.

The proof can readily be outlined by reference to a specific example. Let us assume that p^* has been chosen as .999 (and similarly, that ϵ, for use of the rule (11.1), has been chosen as .001). Then consider the case, mentioned earlier, of ten successive flippings of a given coin. Choose as "premise" the statement e: 'The statistical probabilities of getting heads and of getting tails by flipping this coin are both ½; the results of different flippings are statistically independent; and S is a particular sequence of 10 flippings of this coin'; furthermore, let h_1 be 'S does not yield tails 10 times in succession'; h_2: 'S does not yield 9 tails followed by 1 head'; h_3: 'S does not yield 8 tails followed by 2 heads'; and so on to h_{1024}: 'S does not yield heads 10 times in succession.' Each of these hypotheses h_j ascribes to S a certain kind of outcome O_j; and as is readily seen, the probability statements included in e imply logically that for each of these 2^{10} different possible outcomes, the statistical probability of obtaining it as a result of 10 successive flippings of the given coin is $1 - (½)^{10}$. But according to rule (11.1), this makes it practically certain, for any one of the O_j that the particular sequence S will have O_j as its outcome; in other words, this makes it practically certain, for each one of our hypotheses h_j, that h_j is true. Rule (11.2) more specifically ascribes the logical probability $1 - (½)^{10}$ to each of the h_j on the basis of the statistical probability for O_j which is implied by e.[61]

On the other hand, the conjunction of the h_j is tantamount to the assertion that none of the 10 particular flippings that constitute the individual sequence S will yield either heads or tails—a kind of outcome, say O^*, for which e implies the statistical probability zero. This, under rule (11.1), makes it practically certain that this outcome will not occur in S, i.e., that the conjunction of the h_j is false—even though each of the conjoined hypotheses is practically certain to be true. And if (11.2) is invoked, then the statement that the statistical probability of O^* is zero confers upon the conjunction of the h_j the logical probability zero even though, on the basis of statistical information also provided by e, each of the h_j has a logical probability exceeding .999.[62]

A similar argument can be presented for the case, considered earlier, of the radioactive decay of a particular sample S of 10 milligrams of radon over a period of 7.64 days. For the interval from 2 to 3 milligrams referred to in our previous discussion can be exhaustively divided into mutually exclusive subintervals $i_1, i_2, \ldots, i_n$, which are so small that for each i_j there is a statistical probability exceeding .999999, let us say, that the residual mass of radon left of an initial 10 milligrams after 7.64 days will *not* lie within i_j. Hence, given the information that the half life of radon is 3.82 days, it will be practically certain, according to rule (11.1) that if the experiment is performed just with the one particular sample S, the residual mass of radon will not lie within the interval i_1; it will also be practically certain that the residual mass will not lie within i_2; and so forth. But conjointly these hypotheses, each of which is qualified as practically certain, assert that the residual mass will not lie within the interval from 2 to 3 milligrams; and as was noted earlier, the law stating the half life of radon makes it practically certain that precisely the contradictory of this assertion is true! Thus, the statistical information about the half life of radon statistically explains (or predicts, etc., depending on the context) to a very high degree each of the individual hypotheses referring to the subintervals; but it does not thus explain (or predict, etc.) their conjunction.

Though superficially reminiscent of the ambiguity of statistical systematization, which was examined earlier, this nonadditivity is a logically quite different characteristic of statistical systematization. In reference to statistical systematizations of the simple kind suggested by rule (11.1), ambiguity can be characterized as follows: If the fact that b is G can be sta-

[61] Thus the basis for the assignment, under rule (11.1), of the probability $1 - (½)^{10}$ to each h_j is, strictly speaking, not e, but the sentence e_j: 'The statistical probability of obtaining O_j as a result of 10 successive flippings of the coin is $1 - (½)^{10}$, and S is a particular set of n such successive flippings'; this e_j is a logical consequence of, but not equivalent to, e. Now, in general, if $c(h^*, e^*) = q$ and e^{**} is a logical consequence of e^*, then $c(h^*, e^{**})$ need not equal q at all; but in our case, it is extremely plausible to assume that whatever information e contains beyond e_j is inductively irrelevant to h_j; and on this assumption, we then have $c(h_j, e) = 1 - (½)^{10}$ for each j. The requisite assumption may also be expressed more generally in the following rule, which is a plausible extension of (11.2):

Let e be a sentence which (i) specifies, for various outcomes G_k of a random experiment F, their statistical probabilities $p(G_k, F) = r_k$, (ii) states that the outcomes of different performances of F are statistically independent of each other, and (iii) asserts that a certain particular event b is a case of n successive performances of F; and let e assert nothing else. Next, let h be a statement ascribing to each of the particular performances of F that constitute b some particular one of the various outcomes G_k. Then $c(h, e)$ equals the product of the statistical probabilities of those n outcomes. (For example, if b consists of three performances of F and h asserts that the first and third of these yield G_2, and the third G_4, then $c(h, e) = r_2 \cdot r_4 \cdot r_2$.)

[62] The observation made in the preceding note applies here in an analogous manner.

tistically explained (predicted) by a true explanans stating that b is F and that $p(G, F) > 1 - \epsilon$, then there is in general another true statement to the effect that b is F′ and that $p(-G, F') > 1 - \epsilon$, which in the same sense statistically explains (predicts) that b is non-G. This ambiguity can be prevented by requiring that a statistical systematization, to be scientifically acceptable, must satisfy the principle of total evidence; for one and the same body of evidence cannot highly confirm both 'Gb' and '—Gb.'

But the principle of total evidence does not affect at all the nonconjunctiveness of statistical systematization, which lies precisely in the fact that one and the same set of inductive "premises" (one and the same body of evidence) e may confirm to within $1 - \epsilon$ each of n alternative "conclusions" (hypotheses), while confirming with equal strength also the negation of their conjunction. This fact is rooted in the general multiplication theorem of elementary probability theory, which implies that the probability of the conjunction of two items (characteristics or statements, according as statistical or logical probabilities are concerned) is, in general, less than the probability of either of the items taken alone. Hence, once the connection between "premises" and "conclusion" in a statistical systematization is viewed as probabilistic in character, nonconjunctiveness appears as inevitable, and as one of the fundamental characteristics that distinguish statistical systematization from its deductive-nomological counterpart.

14. *Concluding Remarks.*

Commenting on the changes that the notion of causality has undergone as a result of the transition from deterministic to statistical forms of physical theory, R. von Mises holds that "people will gradually come to be satisfied by causal statements of this kind: It is *because* the die was loaded that the 'six' shows more frequently (but we do not know what the next number will be); or, *Because* the vacuum was heightened and the voltage increased, the radiation became more intense (but we do not know the precise number of scintillations that will occur in the next minute)."[63] This passage clearly refers to statistical explanation in the sense considered in the present essay; it sets forth what might be called a statistical-probabilistic concept of "because," in contradistinction to a strictly deterministic one, which would correspond to deductive-nomological explanation. Each of the two concepts refers to a certain kind of subsumption under laws—statistical in one case, strictly universal in the other; but, as has been argued in the second part of this study, they differ in a number of fundamental logical characteristics: The deterministic "because" is deductive in character, the statistical one is inductive; the deterministic "because" is an either-or relation, the statistical one permits degrees; the deterministic "because" is unambiguous, while the statistical one exhibits an ambiguity which calls for relativization with respect to the total evidence available; and finally, the deterministic "because" is conjunctive whereas the statistical one is not.

The establishment of these fundamental logical differences is at best just a small contribution toward a general analytic theory of statistical modes of explanation and prediction. The fuller development of such a theory raises a variety of other important issues, some of which have been touched upon in these pages; it is hoped that those issues will be further clarified by future investigations.

[63] Mises [34], p. 188; italics in original text.

REFERENCES

1. Alexander, H. Gavin. "General Statements as Rules of Inference?" in *Minnesota Studies in the Philosophy of Science*, Vol. II, H. Feigl, M. Scriven, and G. Maxwell, eds. Minneapolis: University of Minnesota Press, 1958. Pp. 309–329.
2. Barker, S. F. *Induction and Hypothesis.* Ithaca: Cornell University Press, 1957.
3. Braithwaite, R. B. *Scientific Explanation.* Cambridge: Cambridge University Press, 1953.
4. Campbell, Norman. *What Is Science?* New York: Dover Press, 1952.
5. Carnap, R. *The Logical Syntax of Language.* New York: Harcourt, Brace, and Co., 1937.
6. Carnap, R. "On Inductive Logic," *Philosophy of Science*, 12:72–97 (1945).
7. Carnap, R. *Logical Foundations of Probability.* Chicago: University of Chicago Press, 1950.
8. Carnap, R. *The Continuum of Inductive Methods.* Chicago: University of Chicago Press, 1952.
9. Carnap, R., and Y. Bar-Hillel. *An Outline of a Theory of Semantic Information.* Massachusetts Institute of Technology, Research Laboratory of Electronics. Technical Report No. 247. 1952.
10. Cohen, M. R., and E. Nagel. *An Introduction to Logic and Scientific Method.* New York: Harcourt, Brace, and Co., 1934.
11. Cramér, H. *Mathematical Methods of Statistics.* Princeton: Princeton University Press, 1946.
12. De Finetti, Bruno. "Recent Suggestions for the Reconciliations of Theories of Probability," in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, J. Neyman, ed. Berkeley: University of California Press, 1951. Pp. 217–226.
13. Dray, W. *Laws and Explanation in History.* London: Oxford University Press, 1957.
14. Duhem, Pierre. *La Théorie physique, son objet et sa structure.* Paris: Chevalier et Rivière, 1906.

15. Feigl, H., and May Brodbeck, eds. *Readings in the Philosophy of Science*. New York: Appleton-Century-Crofts, 1953.
16. Feigl, H., and W. Sellars, eds. *Readings in Philosophical Analysis*. New York: Appleton-Century-Crofts, 1949.
17. Feigl, H., M. Scriven, and G. Maxwell, eds. *Minnesota Studies in the Philosophy of Science*, Vol. II. Minneapolis: University of Minnesota Press, 1958.
18. Galilei, Galileo. *Dialogues Concerning Two New Sciences*. Transl. by H. Crew and A. de Salvio. Evanston, Ill.: Northwestern University, 1946.
19. Gardiner, Patrick, ed. *Theories of History*. Glencoe, Ill.: Free Press, 1959.
20. Goodman, Nelson. "The Problem of Counterfactual Conditionals," *Journal of Philosophy*, 44:113–128 (1947). Reprinted, with minor changes, as the first chapter of Goodman [20].
21. Goodman, Nelson. *Fact, Fiction, and Forecast*. Cambridge, Mass.: Harvard University Press, 1955.
22. Goodman, Nelson. "Recent Developments in the Theory of Simplicity," *Philosophy and Phenomenological Research*, 19:429–446 (1959).
23. Hanson, N. R. "On the Symmetry between Explanation and Prediction," *Philosophical Review*, 68:349–358 (1959).
24. Hempel, C. G. "The Function of General Laws in History," *Journal of Philosophy*, 39:35–48 (1942). Reprinted in Feigl and Sellars [16], and in Jarrett and McMurrin [27].
25. Hempel, C. G. "The Theoretician's Dilemma," in *Minnesota Studies in the Philosophy of Science*, Vol. II, H. Feigl, M. Scriven, and G. Maxwell, eds. Minneapolis: University of Minnesota Press, 1958. Pp. 37–98.
26. Hempel, C. G., and P. Oppenheim. "Studies in the Logic of Explanation," *Philosophy of Science*, 15:135–175 (1948). Secs. 1–7 of this article are reprinted in Feigl and Brodbeck [15].
27. Jarrett, J. L., and S. M. McMurrin, eds. *Contemporary Philosophy*. New York: Henry Holt, 1954.
28. Jeffrey, R. C. "Valuation and Acceptance of Scientific Hypotheses," *Philosophy of Science*, 23:237–246 (1956).
29. Kemeny, J. G., and P. Oppenheim. "Degree of Factual Support," *Philosophy of Science*, 19:307–324 (1952).
30. Körner, S., ed. *Observation and Interpretation*. Proceedings of the Ninth Symposium of the Colston Research Society. New York: Academic Press Inc., 1957. London: Butterworth, 1957.
31. Kolmogoroff, A. *Grundbegriffe der Wahrscheinlichkeitrechnung*. Berlin: Springer, 1933.
32. Lewis, C. I. *An Analysis of Knowledge and Valuation*. La Salle, Ill.: Open Court Publishing Co., 1946.
33. Luce, R. Duncan, and Howard Raiffa. *Games and Decisions*. New York: Wiley, 1957.
34. Mises, Richard von. *Positivism. A Study in Human Understanding*. Cambridge, Mass.: Harvard University Press, 1951.
35. Nagel, E. *Logic without Metaphysics*. Glencoe, Ill.: The Free Press, 1956.
36. Neumann, John von, and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press, 2d ed., 1947.
37. Neyman, J. *First Course in Probability and Statistics*. New York: Henry Holt, 1950.
38. Popper, K. R. *Logik der Forschung*. Vienna: Springer, 1935.
39. Popper, K. R. "The Propensity Interpretation of the Calculus of Probability, and the Quantum Theory," in *Observation and Interpretation*, S. Körner, ed. Proceedings of the Ninth Symposium of the Colston Research Society. New York: Academic Press Inc., 1957. London: Butterworth, 1957. Pp. 65–70.

40. Popper, K. R. *The Logic of Scientific Discovery*. London: Hutchinson, 1959.
41. Reichenbach, H. *The Theory of Probability*. Berkeley and Los Angeles: University of California Press, 1949.
42. Rescher, N. "On Prediction and Explanation," *British Journal for the Philosophy of Science*, 8:281–290 (1958).
43. Rescher, N. "A Theory of Evidence," *Philosophy of Science*, 25:83–94 (1958).
44. Ryle, G. *The Concept of Mind*. London: Hutchinson, 1949.
45. Savage, L. J. *The Foundations of Statistics*. New York: Wiley, 1954.
46. Scheffler, I. "Prospects of a Modest Empiricism," *Review of Metaphysics*, 10: 383–400, 602–625 (1957).
47. Scheffler, I. "Explanation, Prediction, and Abstraction," *British Journal for the Philosophy of Science*, 7:293–309 (1957).
48. Schlick, M. "Die Kausalität in der gegenwärtigen Physik," *Die Naturwissenschaften*, 19:145–162 (1931).
49. Scriven, M. "Definitions, Explanations, and Theories," in *Minnesota Studies in the Philosophy of Science*, Vol. II, H. Feigl, M. Scriven, and G. Maxwell, eds. Minneapolis: University of Minnesota Press, 1958. Pp. 99–195.
50. Scriven, M. "Explanations, Predictions, and Laws," in this volume of *Minnesota Studies in the Philosophy of Science*, pp. 170–230.
51. Sellars, W. "Inference and Meaning," *Mind*, 62:313–338 (1953).
52. Sellars, W. "Conterfactuals, Dispositions, and the Causal Modalities," in *Minnesota Studies in the Philosophy of Science*, Vol. II, H. Feigl, M. Scriven, and G. Maxwell, eds. Minneapolis: University of Minnesota Press, 1958. Pp. 225–308.
53. Toulmin, S. *The Philosophy of Science*. London: Hutchinson, 1953.
54. Wald, A. *Statistical Decision Functions*. New York: Wiley, 1950.
55. Williams, D. C. *The Ground of Induction*. Cambridge, Mass.: Harvard University Press, 1947.