Introduction

Suppose that a small plane crashed upon takeoff from an airport near Denver on July 15, 1971. We ask why the crash occurred. Our interest in the event might be highly practical; the FAA, for instance, investigates such accidents in order to improve flying safety. When we know why occurrences of various types happen, we can often do something about controlling them. At the same time, our interests might be largely theoretical. To someone concerned with aerodynamics, the search for an explanation of the crash might be the result of sheer intellectual curiosity. In either case—these two motives are not mutually exclusive—we look to science for an explanation of the event; the explanation may be both practically useful and intellectually satisfying. Whether we pursue scientific investigations for the purpose of predicting and controlling our environment or for the sake of understanding the world in which we live, the search for explanations is at the heart of the endeavor.

As the inquiry gets underway, the investigators will establish a number of relevant facts, such as the type of aircraft involved, its mechanical condition, the load it was carrying, the length of the runway, and the height and location of the obstacle. In addition, they will take into account such relevant meteorological circumstances as the wind velocity, the atmospheric pressure, the temperature, and the relative humidity. Given all of these conditions, the investigators can determine the distance needed for takeoff to clear the offending barrier. Having ascertained that the wind was calm and having ruled out such causes as mechanical failure, they find that the atmospheric pressure was low (due to the high altitude of the airport) and that the day was hot and humid. Since the distance required for takeoff depends upon the density of the air—the smaller the density the greater the distance needed—and since density decreases as altitude, temperature, and humidity are increased, the conditions at the time and place of the accident resulted in an abnormally long takeoff distance. Under these circumstances the runway simply was not long enough. The pilot made the fatal error of failing to take these factors into account.

This example illustrates several important features of scientific explanation. We have a particular event—the crash of the airplane—for which an explanation is sought; it is known as the *explanandum* (that which is to be explained).[1] It is explained by invoking *general laws*—for example, that there is an inverse relation between the density of the air and the distance needed for takeoff—under which the explanandum can be subsumed. The explanandum is brought under the general laws by establishing such *initial conditions* as air temperature, atmospheric pressure, relative humidity, wind velocity, type of aircraft, and height and location of the obstacle. The general laws and the initial conditions together constitute the *explanans* (that which does the explaining). The explanatory facts which make up the explanans thus consist of *particular facts*—embodied in the initial conditions—and *general facts*—embodied in the general laws. In order to explain a particular occurrence, both types of facts are essential. The fact that the relative humidity was high on the day of the crash will not help to explain the crash unless we have a general law relating humidity to air density. Similarly, the general relationship between humidity and density is useless without the particular value of the humidity at the place and time of the crash, which is required to bring the generalization to bear on that particular occurrence. General laws are needed to relate particular explanatory facts to the explanandum; particular facts are needed to make the general laws applicable to the explanandum. For obvious reasons, explanations that conform to this pattern are called *covering-law explanations,* and the pattern itself is known as the *covering-law model* of explanation.

We have not, of course, offered a full explanation of the crash; in order to do so, it would be necessary to fill in many details that we have only sketched. Even if such details had been furnished, thereby providing a complete explanation of the crash, it might still seem reasonable to ask for explanations of one or more parts of the explanans. This does not mean that the explanation of the crash is incomplete, but only that there are *other* explanations that might be in order. For example, it may be fairly evident why air is less dense at greater altitudes than at lesser ones, but perhaps it is puzzling that humid air is less dense than dry air. This *general fact* can be explained by noting that, at specified values of pressure and temperature, a particular volume of gas contains approximately the same number of molecules regardless of the kinds of molecules composing it (Avogadro's law). Dry air contains mostly nitrogen (N_2) and oxygen (O_2) molecules, whereas humid air contains a significant proportion of water (H_2O) molecules. The molecular weights of N_2 and O_2 are 28 and 32 respectively, whereas the molecular weight of water vapor is 18. Although a wet washcloth is obviously much heavier than a dry one, a given volume of humid air is less massive, and hence less dense, than the same volume of dry air (at the same pressure and temperature).

We see, then, that general laws, as well as particular facts, are amenable to scientific explanation, and that the general law is explained by subsuming it under still broader laws. Thus, the general relation between the density of moist and dry air is explained in terms of still more general laws relating the density of a gas to its molecular composition. In this book we shall be concerned mainly with explanations of particular events rather than general laws, but it is important to remember that the general laws employed in such explanations are, themselves, capable of being explained by means of covering-law explanations.

In the context of aeronautical engineering, it is appropriate to regard the laws of physics as strict universal generalizations that hold without exception. Thus, the Bernoulli principle, which determines the lift of a wing, can be taken as an unexceptionable law relating the velocity of flow of a fluid (liquid or gas) to the pressure it exerts in a direction perpendicular to the direction of flow. From a more precise and theoretical standpoint, however, we must regard such laws as statistical generalizations that admit of overwhelmingly improbable exceptions. According to this more refined conception, the performance of an airplane attempting a takeoff is determined by the average behavior of exceedingly large numbers of molecules that collide with the propeller, wings, control surfaces, and other parts of the craft. If, for instance, an extremely large number of molecules of air near the obstacle, in the course of their purely random motions, had chanced to be moving upward at just the proper moment, they could have lifted the airplane over the obstacle, thus avoiding the accident. Such phenomena are actually observed for microscopic particles in Brownian motion, but for an object the size of an airplane such an occurrence is so incredibly improbable that for all practical purposes we can ignore its possibility. An occurrence of this type would be analogous to Jeffrey's example of a tire inflating spontaneously.[2]

The fact that certain laws are statistical in character, rather than strictly universal, obviously does not preclude their use in scientific explanations. If we ask why a particular ice cube melted, it would be adequate to point out that it was placed in lukewarm water and that an ice cube in these circumstances will very, *very* probably absorb heat from the surrounding water. It is not physically impossible, however, for the ice cube to give up some of its heat to the water, increasing rather than decreasing the temperature difference between them. Similarly, in strictest rigor the FAA investigators ought to say that an airplane of specified characteristics will very probably follow a particular flight path under particular meteorological conditions. This probabilistic construal of the explanation does not deprive it of its explanatory power. It is still a covering-law type of explanation, but the law is statistical rather than universal.[3]

It is evident from the foregoing remarks that statistical generalizations can be explained in much the same fashion as universal generalizations, namely, by subsumption under broader generalizations. Thus, for instance, the fact that two bodies of unequal temperature will probably exchange heat when brought into thermal contact, the warmer losing heat to the cooler until a uniform temperature is attained throughout, is explained by statistical generalizations concerning the exchange of kinetic energy among colliding molecules. In similar fashion, an explanation of the Bernoulli principle *as a statistical generalization* can also be grounded in the theory of the statistical behavior of gas molecules, inasmuch as pressure exerted on an airfoil by a gas is understood in terms of the averages of vast numbers of collisions of gas molecules with the surface.

The covering-law model of explanation demands, as we have seen, that every explanans contain at least one general law. The laws invoked in the explanation may be either universal or statistical generalizations. Whichever type of law is employed in the explanans, there seems to be a hierarchy of explanations, beginning at the lowest level with explanations of particular events and progressing upward through explanations of general laws of greater and greater scope. In fact, it has often been suggested that particular observable events, such as the airplane crash, can be explained through the use of *empirical laws*, whereas empirical laws are in turn explained by means of *theories*. This distinction between empirical laws and theories depends upon a rough distinction between observables and unobservables. An empirical law embodies general relations among more or less directly observable things and their more or less directly observable properties, whereas theories make reference to unobservables. Airplanes, runways, cylinders of air, and barometers, for instance, are observable entities, whereas molecules of nitrogen and oxygen are not. Likewise, the atmospheric pressure, the weight of an airplane, the wind velocity, and the length of the runway are observable properties of observable things, whereas the number of molecules in a container of gas and the kinetic energy of a single molecule are not. Consequently, the inverse relation between relative humidity and takeoff distance is an *empirical law* which can be used to explain a particular event, the airplane crash, but the generalizations which involve weights of individual molecules and numbers of molecules in a given volume of gas are highly theoretical in character. It was by means of such *theoretical laws* that we explained the foregoing *empirical law*, and this seems typical of the way in which empirical laws are theoretically explained. Theoretical laws, themselves, are often explained in terms of more general and fundamental theories. In this book we shall confine attention chiefly to the explanations of particular events which stand at the bottom of the hierarchy. Once the explanations of the lowest level are understood, we can perhaps hope to move upward and cope with higher-level explanations.

Although the foregoing examples illustrate much of what is involved in scientific explanations, there remains the surprisingly difficult *philosophical* task of providing a general characterization of the logical structure of scientific explanation. This is an ancient problem; as Jeffrey points out in his essay, it goes back at least to Aristotle. Nevertheless, it has been the object of vigorous and fruitful investigation for the past quarter century, due largely to the work of Carl G. Hempel. Beginning with a classic article in 1948, Hempel has elaborated a remarkably appealing and admirably precise theory of the nature of scientific explanation.[4] Although Hempel's account has not been universally accepted by philosophers of science, it has been extremely influential, and it constitutes the closest thing we have to a received view. An extensive literature has grown up around it. For these reasons Hempel's account serves as the very best point of departure for any contemporary discussion of scientific explanation.

Hempel has advanced two basic models of scientific explanation of particular events. The first is the so-called *deductive-nomological* (D-N) pattern; its explanans consists of *universal* laws and initial conditions, from which the occurrence of the explanandum follows deductively. The deductive-nomological explanation *is* a valid *deductive argument;* the statements asserting the initial conditions and universal laws of the explanans are the premises, and the assertion of the explanandum is the conclusion. The second is the so-called *inductive-statistical* (I-S) pattern. The general law in this type of explanation is a *statistical* generalization, and the explanation *is* an *inductive argument.* The statements of initial conditions and statistical laws in the explanans do not necessitate the explanandum, but they do confer high inductive probability upon it. Harking back to our example of the airplane crash for illustration, we can see that it would be possible in principle, though enormously complicated in fact, to construct a valid deduction of the occurrence from strictly universal laws and sufficiently detailed specifications of initial conditions. If, however, we decide for the sake of greater rigor to treat the laws of aerodynamics as statistical generalizations, we could presumably construct an inductive argument according to which the crash was overwhelmingly probable. The logical details of these two models of scientific explanation are presented, along with simple examples of their application, in section 1 of my essay, "Statistical Explanation." For the moment, it is sufficient to observe that *explanations of each type are arguments,* deductive or inductive, *showing that the event to be explained*—the explanandum—*was to be expected by virtue of the explanatory facts set forth in the explanans.*

Deductive validity is an all-or-nothing affair; deductive arguments are either valid or invalid, and there are no degrees of validity. Inductive support, by contrast, does admit of degrees. The premises of an inductive argument may lend more or less weight to the conclusion, and one may speak of degrees of strength of inductive inferences. According to Hempel's view, an inductive-statistical explanation, therefore, has a degree of strength which he designates as the *inductive probability* conferred upon the explanandum by the explanans.

Deductive logic is a highly developed discipline, and any questions about deductive validity that are likely to arise in the course of the discussion of scientific explanation can be settled rather straightforwardly. Inductive logic, by contrast, is in a rather primitive state of development, and many fundamental problems about inductive-statistical explanation cannot easily be settled by referring to a canonical system of inductive logic. If, for example, we were to ask for further elucidation of the concept of inductive probability, it would be reasonable to go to Rudolf Carnap's theory of confirmation, the most rigorous and extensive system of inductive probability available at present.[5] When we look at Carnap's inductive logic, we discover a shocking fact: in that system of inductive logic (the one to which Hempel explicitly refers in connection with the concept of inductive probability), *there is no such thing as inductive inference* in the sense required for Hempel's account of inductive-statistical explanation! In Carnap's inductive logic there are no inductive arguments consisting of premises and conclusion, which allow you to affirm the conclusion (with some degree of probability) if you are prepared to assert the premises. On this view, inductive logic is strongly disanalogous to deductive logic, even to the extent of proscribing inference entirely.

This is not the place to go into Carnap's reasons for denying the possibility of inductive inferences or to argue the merit of his arguments. Nor do I mean to suggest that Carnap's system of inductive logic is the only possible one. But an important heuristic point is in order. Since Carnap's account of inductive probability is the most prominent and best-developed theory available, there does seem to be good reason to wonder whether statistical explanations are arguments at all. This is precisely the revolutionary move made by Jeffrey in his essay, "Statistical Explanation vs. Statistical Inference," and it is the first step in developing an alternative to Hempel's models. Jeffrey concludes, roughly speaking, that the statistical explanation of an event *exhibits that event* as the result of a stochastic process from which such events arise with some probability whose degree may be high, middling, or even very low. Exhibition of such a process does not constitute an argument at all, let alone an argument to the effect that the explanandum was to be expected by virtue of its high probability.

If we take seriously the suggestion that some events have low probabilities, a further reason emerges for advancing the alternative to Hempel's theory of explanation. Were we to insist, with Hempel, that a statistical explanation must embody a high probability, then events that are intrinsically improbable, even though they sometimes occur, would consequently defy all explanation. For example, in the light of current physical theory, the spontaneous radioactive decay of a uranium atom may be due to an alpha particle tunneling through the potential barrier of the nucleus. As the alpha particle bombards this nuclear wall, there is a probability of the order of 10^{-38} that it will escape, and there are no further relevant factors to determine in which instance it tunnels out. On Hempel's view such low probability events are in principle incapable of being explained, but on the alternative account quantum mechanics provides an explanation by furnishing all of the facts relevant to their occurrence.

There is, of course, a very strong temptation to maintain that there must be *some* reason why the alpha particle gets through on one occasion, when it fails on so many others, but the reasons are not known at present. This view involves an a priori commitment to determinism—that is, to the doctrine that every event that happens is completely determined by previous causes. On this view our reliance upon probabilities is simply a reflection of our ignorance; further investigation will reveal the unknown causes and enable us to give a full (deductive-nomological) explanation of the event in question. This position seems to me untenable. I do not mean to argue that present physical theory is complete and correct but, rather, that there is no reason to make an a priori decision as to the nature of further physical theories. Perhaps, in the future, improved theories will provide a deterministic account of events that current theory regards as causally undetermined—but perhaps they will not. We should be prepared for the possibility that the indeterministic character of physical theory is correct and that there are events which are intrinisically improbable, not merely improbable in relation to our present incomplete knowledge. In that case, we need an account of statistical explanation that will characterize the explanation of events in terms of statistical laws. It seems desirable for a theory of explanation to admit the possibility of events that are intrinsically undetermined, indeed, events whose intrinsic probabilities are low, without denying the possibility of explaining them. To deny that any events are undetermined seems to involve an unwarranted a priori commitment to determinism; to say that only events with high probability can be explained involves, among other disadvantages, the acute embarrassment of trying to say in some nonarbitrary way how high is high enough.

Jeffrey, Greeno, and I all agree that statistical explanations need not be regarded as inductive arguments, and we agree that a high probability is not required for a correct statistical explanation. If high probability is not the desideratum, what can we offer as a substitute? The answer is *statistical relevance*. This is the view I have tired to elaborate in detail in my essay "Statistical Explanation." To see why statistical relevance is the key concept, consider the case of a person who experiences relief from a neurotic symptom while (or shortly after) undergoing psychotherapy. Does the psychotherapeutic treatment explain the remission of the symptom? The answer to this question depends not only upon the probability of the abatement of symptoms during (or shortly after) therapy; rather, it depends upon the relation between the remission rate for patients undergoing a particular type of treatment and the spontaneous-remission rate. Even if the probability of the remission of symptoms for patients in psychotherapy were very high, that would have no explanatory value if the spontaneous-remission rate were equally high. At the same time, even if the recovery rate for patients were quite low, but still higher than the spontaneous-remission rate, the fact that the individual had submitted to treatment would have some explanatory force in relation to his psychic improvement.

To say that a certain factor is *statistically relevant* to the occurrence of an event means, roughly, that *it makes a difference to the probability of that occurrence*—that is, the probability of the event is different in the presence of that factor than in its absence. This relation of statistical relevance, and its importance to the concept of statistical explanation, is illustrated by a recent development. In my essay, "Statistical Explanation," I introduced as an example the use of vitamin C as a cure for the common cold. At that time I was unaware of Dr. Linus Pauling's views on the efficacy of vitamin C for that purpose, and I quoted what then seemed fairly reliable evidence that the use of vitamin C is statistically irrelevant to recovery from a cold.[6] If Dr. Pauling is right, the use of vitamin C is relevant to recovery, and my previous factual information was incorrect. Clearly, however, the vital question is not "How probable is recovery from a cold if one takes sufficient vitamin C?" but rather "How does the probability for recovery differ between users and nonusers of vitamin C?"

The foregoing considerations allow us to distinguish quite succinctly between Hempel's view and the alternative. Let us dub the alternative account "the *statistical-relevance* model" or "S-R model" for short. The term "inductive" is deliberately omitted from the title to emphasize that S-R explanations are not arguments or inferences of any sort. The two models can be characterized as follows:

I-S model (Hempel): an explanation is an *argument* that renders the explanandum *highly probable.*

S-R model (Jeffrey, Salmon, Greeno): an explanation is an *assembly of facts statistically relevant* to the explanandum, *regardless of the degree of probability* that results.

It is evident that an explanation can satisfy Hempel's high-probability requirement without satisfying the relevance requirement and that the relevance requirement can be fulfilled in the absence of high probability. The fact that high probability is neither necessary nor sufficient for statistical relevance indicates that the difference between Hempel's I-S model and our S-R model is fundamental. In my essay I attempt to explain in detail how one goes about assembling sets of conditions relevant to the occurrence of an event—indeed, even *complete* sets of relevant conditions—and to offer further justification for characterizing statistical explananation in that way. Furthermore, I offer counterexamples—

such as the man who takes birth control pills and avoids becoming pregnant—to show that even deductive-nomological explanations can fail on account of lack of relevance. Richter's excellent drawing which serves as a frontispiece illustrates the same point. We are led to the suggestion that explanations embodying universal generalizations are simply a limiting case of S-R explanation, subject to the same kinds of relevance requirements.

When we look at the airplane crash from the standpoint of relevance, we may start by asking why a brand X airplane in good mechanical condition carrying a reasonable load should fail to clear an obstacle when similar craft with similar loads had often taken off successfully from runways no longer than this one. Whether we are construing the laws as universal or statistical, we want to find *relevant* conditions to account for the crash. The answer, we find, is the air density at the time and place of the crash. It appears that the pilot had made the all-too-common error of forgetting the relevance of altitude, temperature, and humidity to the distance required for takeoff.

Having argued rather adamantly that events with low probabilities are amenable to scientific explanation, I must confess to a feeling of queasiness in saying that an event is explained when we have shown that according to all relevant factors, its occurrence is overwhelmingly improbable. I am somewhat inclined to attribute this feeling to intuitions that have been well nurtured on more than two decades of exposure to Hempel's very persuasive writings, and to say that we simply have to retrain our intuitions. Greeno has a different, and I suspect better, way of handling this matter. Given that a theory (i.e., a collection of statistical laws) has to explain the occurrence and nonoccurrence of many different types of events, and that factors relevant to the nonoccurrence of the event seem to have a place in the explanation,[7] Greeno suggests that we evaluate the overall explanatory power of a theory to explain all of the kinds of events it purports to explain, rather than attempting to evaluate the goodness of a particular explanation of a particular event. One of the attractive features of Greeno's essay, "Explanation and Information," is that it provides an appealing method of assessing at least one aspect of the explanatory value of a theory. This result is achieved by application of some concepts of information theory.

When we ask what good it is to have an S-R explanation, it is satisfying to be able to say that the invocation of an explanation increases our information. Indeed, information theory even provides a quantitative measure of the amount of increase. Perhaps there are other desiderata for explanatory theories, but increase of information is a nice one. Looking at the measure in some detail, we shall see that the addition of information accrues as a result of the fact that the explanation provides appropriate relevance relations.

Consider a simple example. Suppose that the population of Centerville, U.S.A., is equally divided between Democrats and Republicans. Let us call the partition of the population in terms of political affiliation $\{M\}$ (to be thought of as explanandu$\underline{m}$), and let

$$M_1 = D;\ M_2 = R; \text{ where } \{M\} = \{D, R\}.$$

According to our assumption,

$$p_1 = P(M_1) = P(D) = 1/2;$$
$$p_2 = P(M_2) = P(R) = 1/2.$$

This partition involves the greatest possible degree of uncertainty for a partition into two subclasses, for knowing that a person is a resident of Centerville tells us nothing about whether he is a Republican or a Democrat. In information theory this uncertainty is measured by

$$H(M) = \sum_i -p_i log_2 p_i = 1,$$

where it is sometimes called, with enormous potentiality for confusion, the "information." [8] The bifurcation into two equally probable subsets provides the unit of uncertainty (or information) known as the "bit." Notice that the uncertainty achieves its maximum value of 1 when $p_1 = p_2$, and it drops to its minimum of zero when either p_1 or p_2 assumes the value of 1. If all residents of Centerville were Republicans, there would be no uncertainty whatever about their party affiliation.

Now suppose, moreover, that Centerville is split by a set of railroad tracks that run north to south through the town, so that half of the residents live to the east and half live to the west of the tracks. Here we have another partition; let us call it $\{S\}$ (to be thought of as explanan$\underline{s}$), and let

$$S_1 = E;\ S_2 = W; \text{ where } \{S\} = \{E, W\}.$$

According to our second assumption

$$p'_1 = P(S_1) = P(E) = 1/2;$$
$$p'_2 = P(S_2) = P(W) = 1/2.$$

Again, the uncertainty is maximal for such a partition:

$$H(S) = 1.$$

The aggregate uncertainty of the two partitions is their sum,

$$H(M) + H(S) = 2.$$

The important question about these two partitions concerns their mutual independence. Let us call the unconditional probabilities $P(S_i)$ and $P(M_j)$ given above the "marginal probabilities." Let us then introduce the conditional probabilities $P(M_j,S_i) = p_{ij}$ from the members of the partition $\{S\}$ to the members of the partition $\{M\}$.[9] By definition, the partition $\{M\}$ is independent of the partition $\{S\}$ if the conditional probabilities are equal to the respective marginal probabilities:

$$P(M_j, S_i) = P(M_j) \quad \text{or} \quad p_{ij} = p_j \text{ for every } i, j.$$

Intuitively we want to say that the conditional probabilities contribute no further information if the two partitions are statistically independent of one another, but that they can contribute positive information (reduction of uncertainty) if there is a statistical dependency between them. The general idea is this: if the probability of being a Democrat varies, depending upon the side of the tracks on which the resident lives, then knowledge of his place of residence reduces the uncertainty about his party affiliation. If, however, $P(D,E) = P(D,W) = P(D)$, then knowledge of place of residence does not provide any information relevant to party affiliation.

In information theory, the quantitative measure of reduction of uncertainty—that is, the "information transmitted" by the theory T—is given by

$$I_T = H(M) + H(S) - H(S \times M)$$

where

$$H(S \times M) = \sum_i \sum_j -p_i p_{ij} \log_2 p_i p_{ij}.$$

As Greeno shows,

$$H(S \times M) = H(M) + H(S)$$

whenever the conditional probabilities equal the corresponding marginal probabilities (i.e., $p_{ij} = p_j$). Thus, if the partitions are independent, the reduction in uncertainty, or the increase of information, is zero. This, of course agrees with our intuitions. It also means that a partition $\{S\}$ that is statistically irrelevant to a partition $\{M\}$ cannot have any explanatory value with respect to it.

Suppose, however, that the two partitions are not independent and that place of residence is relevant to party affiliation. In particular, let us say that 3/4 of the people to the east of the tracks are Democrats, whereas 3/4 of those on the west side are Republicans:

$$\begin{aligned}
P(D, E) &= P(M_1, S_1) = p_{11} = \tfrac{3}{4};\\
P(R, E) &= P(M_2, S_1) = p_{12} = \tfrac{1}{4};\\
P(D, W) &= P(M_1, S_2) = p_{21} = \tfrac{1}{4};\\
P(R, W) &= P(M_2, S_2) = p_{22} = \tfrac{3}{4}.
\end{aligned}$$

Then,

$$H(S \times M) = -2(\tfrac{3}{8}\, log_2\, \tfrac{3}{8} + \tfrac{1}{8}\, log_2\, \tfrac{1}{8}) \simeq 1.81$$

and

$$I_T \simeq 0.19.$$

This quantity represents our increase of information by virtue of the conditional probabilities.

Greeno shows that the increase in information is maximal when all of the conditional probabilities are either zero or one. This corresponds to the situation in which deductive-nomological explanation is possible. If, however, the marginal probabilities in the original explanandum partition are also either zero or one, that maximum represents no gain in information. This situation corresponds to the case in which deductive-nomological explanation becomes vacuous through failure of relevance conditions, as in the example of the man who takes birth control pills. These considerations show quite clearly that the measure of explanatory value introduced by Greeno is a statistical-relevance measure. An explanatory theory can, on his view, have explanatory value even though it assigns low probabilities to explanandum events, and it may fail to have explanatory value even if it assigns high probabilities to explanadum events.

In my essay, "Statistical Explanation" (section 6), I approach the increase in information resulting from a relevant partition in a different way. Suppose we select a particular resident of Centerville, and ask for the probability that he is a Democrat. We would be ill-advised to accept the value 1/2—say as the basis for a 50–50 bet— because the reference class of residents is not homogeneous with respect to party affiliation. We should look instead at his residence; if he lives on the east side of the tracks, assign the value 3/4, and if he lives on the west side, the value should be 1/4. Let us look at the matter quantitatively by associating the (true) value 1 with each Democrat and the (true) value 0 with each Republican. If we assign the value 1/2 to each person, the error (deviation from the "true" value) in each case is ± 1/2, and (squaring to make everything positive) the squared error is 1/4 for each individual. Obviously, the mean-squared error for all N residents is 1/4. Suppose, instead, that we assign the value 3/4 to each resident east of the tracks, and the value 1/4 to each resident of the west side. The number of residents on each side is $N/2$; $3/4 \times N/2$ of the east siders are Democrats, while $1/4 \times N/2$ are Republicans. The converse situation obtains on the west side. If we assign the value 3/4 to someone who is a Democrat, the error is 1/4 and the squared error is 1/16. If we assign the value 3/4 to a Republican, the error is 3/4, and the squared error is 9/16. The cumulative squared error for all residents of the east side is

$$3/4 \times N/2 \times 1/16 + 1/4 \times N/2 \times 9/16 = 12N/128.$$

The same cumulative squared error occurs if we assign the value 1/4 to each west-sider. The total cumulative squared error is $24N/128$, and the mean-squared error is 3/16, which is less than 1/4, the mean-squared error that results from ignoring the relevant partition in terms of place of residence. We see once more how the increase in information due to a relevant partition of the reference class translates into a numerical measure—in this case, one that is rather obviously related to predictive success.

Greeno has chosen to explicate S-R explanation in terms of information transmitted, whereas I have chosen to explicate it in terms of the homogeneity of the reference class. As we have seen, each approach leads to a quantitative measure, as well as a qualitative characterization. Qualitatively, the explications seem to coincide, for they agree that the essence of explanation is in the relevance relations expressed by the conditional probabilities that relate the explanans partition to the explanandum partition. Both of these treatments provide a straightforward answer to a previously recalcitrant problem concerning the utility of scientific explanations. On the present view some of the values of an S-R explanation are the increase of information, the decrease in uncertainty, and the increase in predictive success provided by genuine explanations. For starters, they seem to supply a reasonable motivation for looking more closely at the structure of S-R explanation.

NOTES

1. For the moment I am being deliberately ambiguous about the nature of the explanandum—whether it is the event or the statement that the event occurs. This technicality will be treated herein in sec. 1 of my essay "Statistical Explanation." Similar remarks apply to the explanans.