



CHICAGO JOURNALS



A Deductive-Nomological Model of Probabilistic Explanation

Author(s): Peter Railton

Source: *Philosophy of Science*, Vol. 45, No. 2 (Jun., 1978), pp. 206-226

Published by: [The University of Chicago Press](#) on behalf of the [Philosophy of Science Association](#)

Stable URL: <http://www.jstor.org/stable/186817>

Accessed: 03/09/2013 13:09

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press and Philosophy of Science Association are collaborating with JSTOR to digitize, preserve and extend access to *Philosophy of Science*.

<http://www.jstor.org>

A DEDUCTIVE-NOMOLOGICAL MODEL OF PROBABILISTIC EXPLANATION*

PETER RAILTON†

Princeton University

It has been the dominant view that probabilistic explanations of particular facts must be inductive in character. I argue here that this view is mistaken, and that the aim of probabilistic explanation is not to demonstrate that the explanandum fact was nomically expectable, but to give an account of the chance mechanism(s) responsible for it. To this end, a deductive-nomological model of probabilistic explanation is developed and defended. Such a model has application only when the probabilities occurring in covering laws can be interpreted as measures of objective chance, expressing the strength of physical propensities. Unlike inductive models of probabilistic explanation, this deductive model stands in no need of troublesome requirements of maximal specificity or epistemic relativization.

What if some things happen by chance—can they nonetheless be explained? How?

Some things *do* happen by chance, according to the dominant interpretation of our present physical theory, the probabilistic interpretation of quantum mechanics. Nonetheless, they can be explained: by that theory, in virtually the same way as deterministic phenomena—deductive-nomologically. At least, that is what I hope to show in this paper.

Our universe may not be deterministic, but all is not chaos. It is governed by laws of two kinds: probabilistic (such as the laws concerning barrier penetration and certain other quantum phenomena) and non-probabilistic (such as the laws of conservation of mass-energy, charge, momentum, etc.).¹ Were the probabilism of laws of the first sort remediable by suitable elaboration of laws of the second sort,

*Received September, 1977.

†I would like to thank C. G. Hempel, Richard C. Jeffrey, and David Lewis for helpful criticisms of earlier drafts. I am especially indebted to David Lewis for the idea that a propensity interpretation of probability sits best with the account of probabilistic explanation given here. I have greatly benefited from discussions of related matters with Sam Scheffler and David Fair.

¹Let us say rather loosely that a system is deterministic if, for any one instant, its state is physically compatible with only one (not necessarily different) state at each other instant. A system is indeterministic otherwise, but lawfully so if a complete description of its state at some one instant plus all true laws together entail a distribution of probabilities over possible states at later times.

Philosophy of Science, 45 (1978) pp. 206–226.
Copyright © 1978 by the Philosophy of Science Association.

the universe would be deterministic after all, and the problem of explaining chance phenomena would no longer be with us. However, indications are that physical indeterminism is irremediable, and that the universe exhibits not only chances, but lawful chances. I will argue that we come to understand chance phenomena, even when the chance involved is extremely remote, by subsuming them under these irremediably probabilistic laws.

1. Introductory Remarks on Explanation. Do I offer a deductive-nomological (*D-N*) model of probabilistic explanation because I believe that nomic subsumption always explains?—No. There are familiar-enough kinds of non-explanatory *D-N* arguments, for example, those that deduce the explanandum from nomically-related symptoms or after-the-fact conditions alone, citing no causes.

Yet it will not do simply to add to the *D-N* model a requirement that the explanans contain causes whenever the explanandum is a particular fact. First, some particular facts may be explained non-causally, e.g., by subsumption under structural laws such as the Pauli exclusion principle. Second, even where causal explanation is called for, the existence of general, causal laws that cover the explanandum has not always been sufficient for explanation: the search for explanation has also taken the form of a search for mechanisms that underlie these laws. ‘Mechanisms,’ however, is not meant to suggest a parochial attitude toward the nomic connections—deterministic or otherwise—that tie the world together and make explanation possible.

An example may help clarify the notion of mechanism appealed to here. The following *D-N* argument suffices to forecast *that* nasty weather lies ahead, but not to explain *why* this is so:

- (S) The glass is falling.
 Whenever the glass falls the weather turns bad.

The weather will turn bad. ([5], p. 106)

Now nothing works like a barometer for predicting the weather, but nothing like a barometer works for changing it. So it is often maintained that (S) lacks explanatory efficacy because barometers lack the appropriate causal efficacy. The following inference, then, remedies the lack of the first because “it proves that the fact is a fact by citing causes and not mere symptoms” ([5], p. 107):

- (C) The glass is falling.
 Whenever the glass is falling the atmospheric pressure is falling.

Whenever the atmospheric pressure is falling the weather turns bad.

The weather will turn bad. ([5], p. 106)

Yet as explanations go, (C) is also lacking: we remain in the dark as to *why* the weather will turn bad. No connection between cause and effect, no mechanism by which falling atmospheric pressure produces a change for the worse in the weather, has been revealed. I do not doubt that some account of this mechanism exists; my point is that its existence is what makes (C) superior to (S) for explanatory purposes.

(C), if moderated by boundary conditions and put less qualitatively, would supply us the capability to predict *and* control the weather (whenever, as in a laboratory simulator, we can manipulate the atmospheric pressure). While prediction and control may exhaust our practical problems in the natural world, the unsatisfactoriness of (C) shows that explanation is an activity not wholly practical in purpose. The goal of understanding the world is a theoretical goal, and if the world is a machine—a vast arrangement of nomic connections—then our theory ought to give us some insight into the structure and workings of the mechanism, above and beyond the capability of predicting and controlling its outcomes. Until supplemented with an account of the nomic links connecting changes in atmospheric pressure to changes in the weather, (C) will explain but poorly. Knowing enough to subsume an event under the right kind of laws is not, therefore, tantamount to knowing the *how* or *why* of it. As the explanatory inadequacies of successful practical disciplines remind us: explanations must be more than potentially-predictive inferences or law-invoking recipes.

Is the deductive-nomological model of explanation therefore unacceptable?—No, just incomplete. Calling for an account of the mechanism leaves open the nature of that account, and as far as I can see, the model explanations offered in scientific texts are *D-N* when complete, *D-N* sketches when not. What is being urged is that *D-N* explanations making use of true, general, causal laws may legitimately be regarded as unsatisfactory unless we can back them up with an account of the mechanism(s) at work. “An account of the mechanism(s)” is a vague notion, and one obviously admitting of degrees of thoroughness, but I will not have much to say here by way of demystification. If one sees what is lacking in (C)—a characterization, whether sketchy or blow-by-blow, of how it is that declining atmospheric pressure effects the changes we describe as “a worsening of the weather,” i.e., a more or less complete filling-in of the links

in the causal chains—one has the rough idea.

The *D-N* probabilistic explanations to be given below do not explain by giving a deductive argument terminating in the explanandum, for it will be a matter of chance, resisting all but *ex post facto* demonstration. Rather, these explanations subsume a fact in the sense of giving a *D-N* account of the chance mechanism responsible for it, and showing that our theory implies the existence of some physical possibility, however small, that this mechanism will produce the explanandum in the circumstances given. I hope the remarks just made about the importance of revealing mechanisms have eased the way for an account of probabilistic explanation that focuses on the indeterministic mechanisms at work, rather than the “nomic expectability” of the explanandum.

2. Hempel’s Inductive-Statistical Model. For Hempel, a statistical explanation (what is called elsewhere in this paper ‘a probabilistic explanation’) is one that “makes essential use of at least one law or theoretical principle of statistical form” ([3], p. 380). Since Hempel distinguishes between statistical laws and mere statistical generalizations, and asserts that the former apply only where “peculiar, namely probabilistic, modes of connection” exist among the phenomena ([3], p. 377), his characterization permits statistical explanation only of genuinely indeterministic processes.² Were some process to have the appearance of indeterminism owing to arcane workings or uncontrolled initial conditions, then no “peculiar . . . probabilistic” modes of connection would figure essentially in explaining this “pseudo-random” process’s outcomes. Not only would statistical explanation be unnecessary for such a process, it would be impossible: no probabilistic *laws* would govern it.

For example, it has been observed that 99% of all cases of infectious mononucleosis involve lymph-gland swelling. The exceptions might be due to a process that randomly misfires 1% of the time. Or, they might arise from the operation of an unknown deterministic mechanism that works to inhibit swelling whenever a patient begins in a particular initial condition, which as a mere matter of fact is typical of 1% of the population. If initial conditions could be partitioned into two mutually exclusive and jointly exhaustive classes *S* and $-S$, such that all *S*s by law eventually develop swelling, and all $-S$ s do not, the generalization ‘99% of all cases of infectious mononucleosis develop lymph-gland swelling’ would have been shown to be no law, but

²Although there is some difficulty in reconciling all that is said in [3] with this conclusion, Hempel now accepts it (personal communication).

merely a descriptive report of observed relative frequencies. No law, it cannot support a statistical explanation. But discovering it not to be a law is just discovering that statistical explanation is uncalled for, since each case of mononucleosis will have been of type S or type $\neg S$ from the outset.

On the other hand, suppose that no such partition of initial conditions exists. Then the presence or absence of swelling is presumably due to a “peculiar . . . probabilistic” connection between disease and symptom, i.e., a real causal indeterminism with probability .99 in each case to produce swelling. The generalization in question would thus be nomological, creating both the possibility and the necessity of statistical explanation.

Given such genuine statistical laws, how does Hempel claim statistical explanation should proceed? He begins his account by distinguishing two sorts of statistical explanation. The first, *deductive-statistical (D-S)* explanation, involves “the deductive subsumption of a narrow statistical uniformity under more comprehensive ones” ([3], p. 380). The second, he argues, is of a qualitatively different sort:

Ultimately . . . statistical laws are meant to be applied to particular occurrences and to establish explanatory and predictive connections among them. ([3], p. 381)

To make such laws relevant to “particular occurrences,” Hempel believes we must go beyond the reach of deduction, and so he proposes an inductive model of statistical explanation.

Inductive-statistical (I-S) explanation proceeds by adducing statistical laws and associated initial conditions relative to which the explanandum is highly probable. High relative probability is required because, on Hempel’s view, statistical laws become explanatorily relevant to an individual chance event only by giving us a basis upon which to inductively infer its occurrence with “practical certainty.” Yet although an *I-S* explanation shows the explanandum to have been “nomically expectable” relative to the explanans, it does not permit detachment of a conclusion; it is less an inference than the expression of an inferential relationship: the explanandum receives a high degree of epistemic support from the explanans. If, for example, we learn that Jones has contracted infectious mononucleosis, we may infer with practical certainty that he will develop lymph-gland swelling. The same inference serves as an *I-S* explanation of the swelling, should it occur. Should it not occur, we would have no explanation for *this*, on Hempel’s model.

However, further investigation of Jones’ medical history might reveal that he suffered mononucleosis once before, and failed to develop

any swelling. Let us suppose that such individuals have a much higher than normal probability of *not* showing swelling in any later bouts with mononucleosis, say .9 rather than .01. This new law and new information about Jones together permit an inference with practical certainty to the conclusion that he will *not* develop swelling, and thus support a corresponding *I-S* explanation. Relative to these new facts, however, no *I-S* explanation would be available should Jones, improbably, develop swelling. What are we to say now about the previous *I-S* explanation, which had just the opposite result? Hempel would reject it as no longer *maximally specific* relative to what we believe about Jones' case. The requirement of maximal specificity is a complicated affair,³ but the basic idea is that we refer each case to the narrowest class of cases to which our present beliefs assign it in which the explanandum has a characteristically different probability. In Jones' case, the narrower class is clearly the class of those contracting mononucleosis for a second time who failed to develop lymph-gland swelling the first time.

If more information about Jones or new discoveries about mononucleosis turn up, we may be forced to move on to still another explanation. *I-S* explanations must be relativized to our current "epistemic situation," and are subject to change along with it. Hempel notes that this sets off *I-S* explanations from *D-N* and *D-S* explanations in a fundamental way:

. . . *the concept of statistical explanation for particular events is essentially relative to a given knowledge situation as represented by a class K of accepted statements. . . .* [W]e can significantly speak of true *D-N* and *D-S* explanations: they are those potential *D-N* and *D-S* explanations whose premises (and hence also conclusions) are true—no matter whether this happens to be known or believed, and thus no matter whether the premises are included in *K*. But this idea has no significant analogue for *I-S* explanation ([3], pp. 402–403)

On Hempel's view, neither of the two contradictory explanations concerning Jones contains false premises, and the explananda in each case do indeed receive the degree of support indicated. It is just that we no longer regard the evidential relationship expressed by the first as explanatorily relevant. Were Jones to develop swelling after all, it would now have to be regarded as inexplicable.

What I take to be the two most bothersome features of *I-S* arguments as models for statistical explanation—the requirement of high proba-

³See, for example, [4].

bility and the explicit relativization to our present epistemic situation (bringing with it an exclusion of questions about the truth of *I-S* explanations)—derive from the inductive character of such inferences, not from the nature of statistical explanation itself. If a non-inductive model for the statistical explanation of particular facts is given, there need be no temptation to require high probability or exclude truth.

3. Jeffrey's Criticism of *I-S* Explanation. Richard C. Jeffrey has criticized Hempel's account on the grounds that statistical explanation is not a form of inference at all, except when the probability of the explanandum is "so high as to allow us to reason, in *any* decision problem, as if its probability were 1" ([5], p. 105). For such exceptional, "beautiful" cases, Jeffrey accepts *I-S* inferences as explanatory because they provide virtual "proof that the phenomenon *does* take place" ([5], p. 106).

For unbeautiful cases, there is no way of proving (in advance) that the explanandum phenomenon will occur. According to Jeffrey, the explanation *why* such unbeatities come to be is a curt "By chance." He has more to say on *how* they come about:

. . . in the statistical case I find it strained to speak of knowledge *why* the outcome is such-and-such. I would rather speak of *understanding the process*, for the explanation is the same no matter what the outcome: it consists of a statement that the process is a stochastic one, following such-and-such a law.⁴ ([5], p. 24)

Jeffrey is surely right, as against Hempel, that probable and improbable outcomes of indeterministic processes are equally explicable, and explicable in the same way. After all, why should it be explicable that a genuinely random wheel of fortune with 99 red stops and 1 black stop came to a halt on red, but inexplicable that it halted on black? Worse, on Hempel's view, halting at any *particular* stop would be inexplicable, even though the wheel must halt at some particular stop in order to yield the explicable outcome *red*.

But I fail to see how Jeffrey can defend his exemption of beautiful cases against a similar line of argument. If the burden in statistical explanation really lies with "*understanding the process* . . . no matter what the outcome," then why should it matter whether the outcome is so highly probable "as to allow us to reason, in *any* decision problem, as if its probability were 1?" The neglect Jeffrey shows here toward minute chances is appropriate for the practical task of decision-making (and perhaps explained by his generally subjectivist approach to

⁴A typographical error has been corrected.

probability), but we must not overlook them in the not-entirely-practical task of explaining. Virtually impossible events may occur, and they deserve and can receive the same explanation as the merely improbable or the virtually certain.

4. A *D-N* Model of Probabilistic Explanation. I will present my account of probabilistic explanation by developing an example of just such “practically negligible”—but physically real and lawful—chance: alpha-decay in long-lived radioactive elements. The mean-life of the more stable radionuclides is so long as to make the probability for any particular nucleus of such an element to decay during our lifetimes effectively zero. But our nuclear theory shows that it is *not* zero, and explains how such rarities can occur.

On the account offered here, probabilistic explanations will be either true or false independent of our epistemic situation. Moreover, to explain, they must be true. Here I am following Hempel’s usage in calling an explanatory argument *true* just in case it is valid and its premises are true ([3], p. 338). Such an explanation will *not* be true if the probabilistic laws it invokes are not true; in particular, it will not be true unless the process responsible for the explanandum is genuinely indeterministic. If alpha-decay is to serve as our paradigm for probabilistic explanation, we must be correct in assuming that the probabilistic wave-mechanical account of particle transmission through the nuclear potential barrier tells us all there is to know about the cause of alpha-decay. At least, it must be true that there are no hidden variables characterizing unknown initial conditions that suffice to account for alpha-decay deterministically. However, I take it to be uncontroversial that alpha-decay is an indeterministic process, if any is.

Let us suppose that we are given an individual instance of alpha-decay to explain: a nucleus of radionuclide Uranium²³⁸, call it ‘*u*’, has emitted an alpha-particle during the time interval lasting from t_0 to $t_0 + \theta$, where θ is very small and expressed in standard units. Since the mean-life of U^{238} is 6.5×10^9 years, the probability of observing a decay by *u* during this interval is exceedingly small, but unquestionably exists (witness the decay). This probability can be given precisely by using the radioactive decay constant λ_{238} characteristic of all atoms of U^{238} . Significantly, we need not know when in the course of the history of *u* time t_0 occurs: the probability of decay is unaffected by the age of the atom. Therefore, as long as decay has not yet occurred, individual “trials”—consisting of observing a single isolated radioactive nucleus for successive intervals of the same length—are statistically independent. Using these two facts we

can determine the probability of decay for individual nuclei during any time interval chosen: it will be 1 minus the probability that any such nucleus *survives* the interval intact; for u , $(1 - \exp(-\lambda_{238} \cdot \theta))$.

To obtain experimental confirmation of this value, we infer *from* the probability to decay of individual nuclei *to* statistical features of sample populations of nuclei, e.g., half-life and mean-life. These predicted statistical features are then checked against actual observed relative frequencies in large populations over long intervals. *Physical* probabilities of the sort being considered here are therefore to be contrasted with *statistical* probabilities; the former express the strength of a certain physical possibility for a given system, while the latter reduce to claims about the (limiting) relative frequencies of traits in sample populations. Much well-founded doubt has been expressed about the applicability of statistical probabilities to single cases, but physical probabilities are *located* in the features of the single case. Therefore, we can understand our nuclear theory as implying strictly universal (physical) probability-attributing laws of the form:

- (1) All nuclei of radioelement E have probability $(1 - \exp(-\lambda_E \cdot t))$ to emit an alpha-particle during any time interval of length t , unless subjected to environmental radiation.

Because schema (1) is universal in form, its instances are candidates for law premises in deductive-nomological inferences concerning individual nuclei. Thus, for u :

- (2) (a) All nuclei of U^{238} have probability $(1 - \exp(-\lambda_{238} \cdot \theta))$ to emit an alpha-particle during any interval of length θ , unless subjected to environmental radiation.

- (b) u was a nucleus of U^{238} at time t_0 , and was subjected to no environmental radiation before or during the interval $t_0 - (t_0 + \theta)$.

-
- (c) u had probability $(1 - \exp(-\lambda_{238} \cdot \theta))$ to emit an alpha-particle during the interval $t_0 - (t_0 + \theta)$.

(2), it appears, gives a $D-N$ explanation only of the fact that u had such-and-such a probability to decay during the interval in question, but we should look a bit closer. I submit that (2), when supplemented as follows, is the probabilistic explanation of u 's decay:

- (3) A derivation of (2a) from our theoretical account of the mechanism at work in alpha-decay.

The $D-N$ inference (2).

A parenthetic addendum to the effect that u did alpha-decay during the interval $t_0 - (t_0 + \theta)$.

Am I merely making a virtue of necessity, and saying that since (3) contains all we can say about u 's decay, (3) must explain it? In fact, there is a great deal more we could say about u 's decay. Deliberately left out of (3) are innumerable details about the experimental apparatus (temperature, pressure, location, etc.), about the beliefs and expectations of those monitoring the experiment, and about the epistemic position of the scientific community at the time. These facts are omitted as *explanatorily irrelevant* to u 's decay because they are *causally irrelevant* to the physical possibility for decay that obtained during the interval in question, and to whether or not that possibility was realized.⁵ A full account of these notions of explanatory and causal relevance is not possible here, so instead I will go on to argue that what (3) comprises *is* explanatorily relevant, and explanatory.

I must begin this task with a defense of the nomological status of (2a), and of the legitimacy of treating it as a covering law for u 's decay. The following criterion of nomologicality will be used: a law is a universal truth derivable from our theory without appeal to particular facts. This criterion of course lacks generality (what counts as theory if not the laws themselves?), fails to segregate natural from logical laws, picks out only so-called "universal" (as opposed to "local") laws, and is entirely too vague (how to distinguish "particular facts" from the rest?). But I trust it will do for now. The motive for excluding "particular facts" is that some true, universal statements derivable from our theory *plus* particular facts would not normally be regarded as universal laws, but would at best be "local laws," e.g., 'All *Homo neanderthalensis* live during the late Pleistocene age.'

The generalization in question here, (2a), is derived by solving the Schrödinger wave-equation for an alpha-particle of energy ≈ 4.2 MeV for the potential regions in and around the nucleus of an element with atomic number 92 and atomic weight 238, none of which are "particular facts," plus some simplifying assumptions about the structure of the nucleus and the distinctness of the alpha-particle within it prior to decay. While it is forbidden by classical physics for a low-energy particle like the ≈ 4.2 MeV alpha-particle associated with U^{238} to pass through the 24.2 MeV potential barrier surrounding

⁵Causal relevance is established here via the wave-equation. I do not mean to suggest that *causal* relevance is the only explanatory kind; cf. the mention of structural laws in section 1.

Some such notion of causal relevance appears to lie behind Salmon's "statistical-relevance" model of probabilistic explanation. Yet what matters is whether a factor enters into the probabilities present, not the statistics they produce.

so massive a nucleus, the quantum theory predicts that the probability amplitude for finding such an alpha-particle outside the potential barrier is nonzero. Thus a transmission coefficient for U^{238} alpha-particles is determined, which, given certain simplifying assumptions about the goings-on inside the nucleus, yields the probability that such a particle will tunnel out of the potential well ‘per unit time for one nucleus,’ namely, λ_{238} ([1], p. 175). (2a) thus neither reports a summary of past observations nor expresses a mere statistical uniformity that scattered initial conditions would lead us to anticipate. Instead, it is a law of irreducibly probabilistic form, assigning definite, physically-determined probabilities to individual systems.

It follows that the derivation of conclusions from (2a) by universal instantiation and *modus ponens* is unexceptionable. Were (2a) but a statistical generalization, properly understood as meaning ‘ $(1 - \exp(-\lambda_{238} \cdot \theta))N$ of U^{238} nuclei in samples of sufficiently large size N , on average, decay during the interval $t_0 - (t_0 + \theta)$,’ it could not undergo universal instantiation, and would not permit detachment of a conclusion about the probability obtaining in a single case.

Further, if the wave-equation does indeed tell us all there is to know about the mechanism involved in nuclear barrier penetration, it follows that nothing more can be said to explain why the observed decay of u took place, once we have shown how (2a) is derived from our account of this mechanism, and established that (2) is valid and that (3)’s parenthetic addendum is true.

Still, does (3) explain why the decay took place? It does not explain why the decay *had to* take place, nor does it explain why the decay *could be expected to* take place. And a good thing, too: there is no *had to* or *could be expected to* about the decay to explain—it is not only a chance event, but a very improbable one. (3) does explain why the decay *improbably* took place, which is how it did. (3) accomplishes this by demonstrating that there existed at the time a small but definite physical possibility of decay, and noting that, by chance, this possibility was realized. The derivation of (2a) that begins (3) shows, by assimilating alpha-decay to the chance process of potential barrier tunneling, how this possibility comes to exist. If alpha-decays are chance phenomena of the sort described, then once our theory has achieved all that (3) involves, it has explained them to the hilt, however unsettling this may be to *a priori* intuitions. To insist upon stricter subsumption of the explanandum is not merely to demand what (alas) cannot be, but what decidedly should not be: sufficient reason that one probability rather than another be realized, that is, chances without chance.

Because of the peculiar nature of chance phenomena, it is explanatorily relevant whether the probability in question was realized, even though there is no before-the-fact explanatory *argument*, deductive or inductive, to this conclusion. Indeed, it is the absence of such an argument that makes a place in probabilistic explanation for a parenthetical addendum concerning whether the possibility became actual in the circumstances given. These addenda may offend those who believe that explanations must always be arguments, but at the most general level explanations are *accounts*, not arguments. It so happens that for deterministic phenomena inferences of a particular kind—*D-N* arguments meeting the desiderata suggested in section 1—are explanatory accounts, and this for good reasons. However, indeterministic phenomena are a different matter, and explanatory accounts of them must be different as well. If the present model is accepted, then almost all of the explanatory burden in probabilistic explanation can be placed on deductive arguments—those characterizing the indeterministic mechanism and those attributing a certain probability to the explanandum. But these arguments leave out a crucial part of the story: did the chance fact obtain?

The parenthetical addendum fills this gap in the account, and communicates information that is relevant to the causal origin of the explanandum by telling us that it came about as the realization of a particular physical possibility. Further, it permits us to chain probabilistic explanations together to make more comprehensive explanations, in which each link is able to bear the full explanatory burden for the fact it covers, and is capable of leading us on to the next fact in the causal sequence being explained. From (2) alone we cannot move directly to an account of what the alpha-particle did to a nearby photographic plate, but only to a probability (and a miserably low one) that this account will be true. The parenthetical addendum to (3) furnishes a non-probabilistic premise from which to begin an account of the condition of the photographic plate: the occurrence of an alpha-decay in the vicinity. Dropping off the addendum leaves an explanation, but it is a *D-N* explanation of the occurrence of a particular probability, not a probabilistic explanation of the occurrence of a particular decay.

The scheme for probabilistic explanation of particular chance facts by nomic subsumption that is being offered here, the *deductive-nomological-probabilistic (D-N-P)* model, is this. First we display (or truthfully claim an ability to display) a derivation from our theory of a law of essentially probabilistic form, complete with an account of how the law applies to the deterministic process in question. The

derived law is of the form:

$$(4a) \quad \forall t \forall x [F_{x,t} \rightarrow \text{Prob}(G)_{x,t} = p]$$

“At any time, anything that is F has probability p to be G .”

Next, we adduce the relevant fact(s) about the case at hand, e :

$$(4b) \quad F_{e,t_0}$$

“ e is F at time t_0 ,”

and draw the obvious conclusion:

$$(4c) \quad \text{Prob}(G)_{e,t_0} = p$$

“ e has probability p to be G at time t_0 .”

To which we add parenthetically, and according to how things turn out:

$$(4d) \quad (G_{e,t_0} / -G_{e,t_0})$$

“(e did/did not become G at t_0).”

Whether a D - N - P explanation is true will depend solely upon the truth-values of its premises and addendum, and the validity of its logic. I leave open what becomes of a D - N - P explanation that contains true laws, initial facts, and addendum, but botches the theoretical account of the laws invoked. Let us simply say that the more botched, the less satisfactory the explanation.

The law premise (4a) will be true if all things at all times satisfy the conditional ‘ $F_{x,t} \rightarrow \text{Prob}(G)_{x,t} = p$ ’, using whatever reading of ‘ \rightarrow ’ we decide upon for the analysis of natural laws in general. It will be false if there exists a partition of the F s into those with *physical* probability r to be G and those with *physical* probability s to be G , where $s \neq r \neq p$. Such a partition might exist according to some *other* interpretation of probability, but this would not affect the truth of (4a). For example, suppose that a coin toss meeting certain specifications is an indeterministic event with probability $\frac{1}{2}$ of yielding heads. We now perform the experiment of repeating such a toss a great many times. Curiously, all and only even numbered tosses yield heads. This result supplies certain frequentists with grounds for saying that $\text{Prob}(\text{heads, even-numbered toss}) = 1$, while $\text{Prob}(\text{heads, odd-numbered toss}) = 0$.

odd-numbered toss) = 0.⁶ But because all tosses met the specification laid down, the probability of heads was the same, $\frac{1}{2}$, on each toss, despite the curious behavior. Such behavior may make us suspicious of our original claims about the indeterminacy of the process or about the physical probability it has of producing heads, but is no proof against them. Indeed, the original probability attribution requires us to assign a definite physical probability to just such an untoward sequence of outcomes, the occurrence of which therefore hardly contradicts this attribution.

The particular fact premise (4b) will be true iff e is an F during the time in question, and not either an F^* (with probability $r \neq p$ to be G) or an F^{**} (with probability $q = p$ to be G , but unlike an F in other respects). Using the (let us say) true law that all F^{**} s have probability $q = p$ to be G , and the falsehood that e is an F^{**} , we could derive a true conclusion, indistinguishable from (4c). Hence the requirement that the *premises* be true if the argument is to explain; and if we reason logically from true premises, the conclusion will take care of itself.

5. Epistemic Relativity and Maximal Specificity Disowned. Have I kept my promise to give an account of probabilistic explanation free from relativization to our present epistemic situation?

Let us return to explanation (3), and admit that it is not the whole story: 23% of the alpha-particles emitted by U^{238} have kinetic energy 4.13 MeV, while the remaining 77% have 4.18 MeV. Therefore there are two different decay constants, $\lambda_{238}^{4.13}$ and $\lambda_{238}^{4.18}$; both are distinct from λ_{238} , used in (3). Hence we must be quite careful in stating what exactly (3) explains. It does *not* explain the particular *event* observed, for this was either a 4.13 or a 4.18 MeV decay, neither of which has probability λ_{238} in unit time. Instead, (3) explains the particular *fact about* the event observed that we set out to explain, namely, that an alpha-decay with unspecified energy (or direction, or angular momentum, etc.) took place at nucleus u during the time interval in question. This fact *does* have probability λ_{238} of obtaining in unit time, representing the sum of the two energy-correlated probabilities with which such a decay might occur.

⁶Cf. the discussion of place selections and homogeneity in [6], sections 4 and 7. Salmon's criterion, which requires formal randomness, would here fail to distinguish a *randomly-produced* regular sequence from a *deterministically-produced* one. Notwithstanding formal similarities, only the latter is appropriately explained non-probabilistically.

If we should learn that the decay of u was of a 4.18 MeV alpha-particle, an explanation of *this* fact would have to be referred to the more specific class of decays with probability $\lambda_{238}^{4.18}$ in unit time. Is the maximal specificity requirement thereby resurrected? There is no need for it. (3) is not an unspecific explanation of this more specific fact, but a fallacious one. It would be logically corrupt to conclude from law (2a) that an individual U^{238} nucleus has probability $(1 - \exp(-\lambda_{238} \cdot \theta))$ to decay *with energy 4.18 MeV* during any interval of length θ , since (2a) says nothing whatsoever about decay energies. The only relevant conclusion to draw from (2a) is (2c), which remains true in the face of our more detailed knowledge about the event in question. Nor is law (2a) falsified by the discovery of a 23:77 proportional distribution of decay energies, and the associated difference in decay rates. For according to our nuclear theory, there is no difference in initial condition between a nucleus about to emit a 4.13 MeV alpha-particle and one about to emit a 4.18 MeV alpha-particle. It remains true that *all* U^{238} nuclei have probability λ_{238} to decay in unit time, but it is further true that all have probability $\lambda_{238}^{4.13}$ to decay one way, and probability $\lambda_{238}^{4.18}$ to decay another.

It must next be determined whether the existence of a difference in probability *due to* a difference in initial condition can be handled by the *D-N-P* model without appeal to a maximal specificity requirement. To permit consideration of possible epistemological complications, it will be assumed that neither the difference in probability nor the partition of initial conditions is known at the start.

Imagine that, although we do not know it, in virtue of certain permanent structural features 23% of all naturally-occurring U^{238} nuclei fall into a class P , and the remaining 77% into a class $-P$, such that only those in P have any probability of emitting a 4.13 MeV alpha-particle, and only those in class $-P$ have any probability of emitting a 4.18 MeV alpha-particle. Suppose further that these two laws have been derived:

- (5) (a) All U^{238} nuclei of type P have probability $(1 - \exp(-\lambda_{238}^{4.13} \cdot t))$ to emit a 4.13 MeV alpha-particle during any time interval of length t , unless subjected to environmental radiation.
- (b) All U^{238} nuclei of type $-P$ have probability $(1 - \exp(-\lambda_{238}^{4.18} \cdot t))$ to emit a 4.18 MeV alpha-particle during any time interval of length t , unless subjected to environmental radiation.

Note that, by our assumptions, the specification of the kinetic energy

of the particle (possibly) emitted may be dropped from (5a) and (5b) without altering the truth of either.

Until the structural differences between types P and $-P$ are discovered and understood, (3) will stand as the accepted explanation of u 's decay. However, once (5a) and (5b) have become known, it will be clear from the fact that u 's alpha-emission had kinetic energy 4.18 MeV that u must have been of type $-P$ prior to decay. Thus a more specific account of u 's decay will be available to scientists, who, already familiar with the theoretical derivation of law (5b), offer the following truncated $D-N-P$ version of this account:

- (6) (a) All nuclei of U^{238} of type $-P$ have probability $(1 - \exp(-\lambda_{238}^{4.18} \cdot \theta))$ to emit an alpha-particle during any time interval of length θ , unless subjected to environmental radiation.
- (b) u was a nucleus of U^{238} of type $-P$ at t_0 , and was subjected to no environmental radiation before or during the interval $t_0 - (t_0 + \theta)$.
-
- (c) u had probability $(1 - \exp(-\lambda_{238}^{4.18} \cdot \theta))$ to emit an alpha-particle during the interval $t_0 - (t_0 + \theta)$.
- (d) (And it did.)

On the Hempelian model (modified so as to permit $I-S$ explanations of improbable phenomena), there is no problem in accounting for the previous acceptability of the $I-S$ counterpart of (3), or for its present unacceptability. (3) had been maximally specific relative to our previous beliefs about alpha-decay in U^{238} , but no longer is, and so is superseded by the more specific (relatively speaking) $I-S$ counterpart of (6).

On the $D-N-P$ model, too, there is no problem in accounting for the acceptability of (3) prior to the discovery of class $-P$ and law (5b): (3)'s premises (and, of course, addendum) were taken to be true. The question is whether, in light of current beliefs, (3) can be ruled out—and (6) ruled in—without invocation of Hempelian constraints. Resolution of the problem (3) and (6) pose through epistemic relativization and maximal specificity requirements seems to me unacceptable. If we were to attribute to nucleus u two unequal probabilities to alpha-decay in a specified way during a single time interval, adding, "Let's pick the most specifically defined value for explanatory purposes," we'd be showing an unseemly tolerance for contradiction in our nuclear theory—and why stop there? Better to face up to the confrontation over truth between (3) and (6), and replace complex and unappealingly relativistic maximal specificity

requirements with the simple requirement of truth. The *D-N-P* model does this. The current unacceptability of (3) is located not in premises insufficiently specific, but in premises insufficiently true, i.e., false. Contrary to (3)'s purported covering law (2a), not all nuclei of U^{238} have probability λ_{238} to decay in unit time if unperturbed by radiation—in fact, none do. In spite of giving accurate expectation values for decay rates in large samples of U^{238} , (2a) is false, and so explanation (3) is ruled out as unsound. Explanation (6), on the other hand, meets the simple requirement of truth, and rules itself in.⁷

Problems about incomplete, misleading, or false beliefs do not bear on whether *D-N-P* explanations have unrelativized truth-values, but concern rather difficulties in *establishing* the truth-values they unrelativistically have. Relativization to our current epistemic situation comes into play only when we begin to discuss whether a given *D-N-P* explanation *seems* true. Whether it *is* true is another matter.

6. Objections to the *D-N-P* Model. I cannot pretend to have said enough about deductive-nomological-probabilistic explanation to have characterized this model adequately. Such reservations as were expressed in section 1 about taking nomic subsumption under a causal law as sufficient for explanation are still in force, and little has been done—except by way of example—to show how the account offered here might accommodate them.

That the probabilistic laws invoked in *D-N-P* explanations are even (in some relevant sense) *causal* cannot be defended until a plausible account of physical probability has been worked out, a task well beyond the scope of this paper. Under a *propensity interpretation*, probability has the characteristics sought: a probability is the expression of the strength of a physical tendency in an individual chance system to produce a particular outcome; it is therefore straightforwardly applicable to single cases; and it is (in a relevant sense) causally responsible for that outcome whenever it is realized. However, propensities are notoriously unclear. For now I can at best assume that clarification is possible, point to a promising start in the attempt to do so—R. N. Giere, “Objective, Single-Case Probabilities and the Foundations of Statistics” ([2])—, and admit that the *D-N-P* model is viable only if sense can be made of propensities, or of objective, physical, lawful, single-case probabilities by any other name.

As for the requirement that explanations elucidate mechanisms, I can only repeat that an essential role is played in *D-N-P* explanations

⁷Explanation (6) is true, however, only under the contrary-to-fact assumption—made for the sake of the example—of the existence of a class $-P$.

by the theoretical deduction of the probabilistic law(s) covering the explanandum.

In lieu of further exposition, I offer the beginnings of a defense, hoping thereby to sketch out the account a bit more fully in those areas most likely to be controversial.

Because it applies only to genuinely indeterministic processes, of which there are few (if any), *D-N-P* explanation is too restricted in scope. It is widely believed that the probabilities associated with standard gambling devices, classical thermodynamics, actuarial tables, weather forecasting, etc., arise not from any underlying physical indeterminism, but from an unknown or uncontrolled scatter of initial conditions. If this is right, then *D-N-P* explanation would be inapplicable to these phenomena even though they are among the most familiar objects of probabilistic explanation. I do not, however, find this troublesome: if something does not happen by chance, it cannot be explained by chance. The use of epistemic or statistical probabilities in connection with such phenomena unquestionably has instrumental value, and should not be given up. What must be given up is the idea that *explanations* can be based on probabilities that have no role in bringing the world's explananda about, but serve only to describe deterministic phenomena.⁸ Whether there *are* any probabilities that enter into the mechanisms of nature is still debated, but the successes of the quantum-mechanical formalism, and the existence of "no hidden variable" results for it, place the burden of proof on those who would insist that physical chance is an illusion.

It could be objected more justly that *D-N-P* explanation is too broad, not too narrow, in scope. Once restrictions have been lifted from the value a chance may have in probabilistic explanation, virtually all explanations of particular facts must become probabilistic. All but the most basic regularities of the universe stand forever in peril of being interrupted or upset by intrusion of the effects of random processes. It might seem a fine explanation for a light's going out that we opened the only circuit connecting it with an electrical power source, but an element of chance was involved: had enough atoms in the vicinity of the light undergone spontaneous beta-decay at the right moment, the electrons emitted could have kept it glowing. The success of a social revolution might appear to be explained by its overwhelming popular support, but this is to overlook the revolu-

⁸Of course, we might speak of statistical or epistemic probabilities as causes of, e.g., beliefs. But if belief formation is not *physically* probabilistic, then probabilistic explanation of it would be impossible, in spite of this sort of causal involvement on the part of statistical or epistemic probabilities.

tionaries' luck: if all the naturally unstable nuclides on earth had commenced spontaneous nuclear fission in rapid succession, the triumph of the people would never have come to pass.

No doubt this proliferation of probabilistic explanations is counter-intuitive, but contemporary science will not let us get away with any other sort of explanation in these cases—it simply cannot supply the requisite non-probabilistic laws. Because they figure in the way things *work*, tiny probabilities appropriately figure in explanations of the way things *are*, even though they scarcely ever show up in the way things turn out.

The *D-N-P* model breaks the link between prediction and explanation. Hempel has justified a “qualified thesis of the structural identity of explanation and prediction” with this principle:

Any rationally acceptable answer to the question ‘Why did *X* occur?’ must offer information which shows that *X* was to be expected—if not definitely, then at least with reasonable certainty. ([3], pp. 367–368)

Abundantly many *D-N-P* explanations—all those covering less than highly probable facts—violate this condition.

However, to abide by this condition and renounce explanations with meager probabilities I take to be worse. Why forgo the explanations of improbable phenomena offered by our theories, when these explanations provide as much of an account of why (and how) their explananda occur as do the explanations of “reasonably certain” phenomena that Hempel’s condition sanctions?

Too restrictive as it stands, Hempel’s condition may be taken in a way not incompatible with *D-N-P* explanation. A *D-N-P* explanation does yield one prediction that is perfectly strict, to the effect that a certain physical probability exists in the circumstances given. If this probability fails to obtain, or to have the value attributed to it, the explanation must be false. It is a complaint against the world, not against the *D-N-P* model, that a direct, non-statistical test for the presence or value of this probability may prove impossible. Remarkably, the mechanisms of the world leave room for spontaneous nuclear disintegrations. Equally remarkably, our physical theory gives us insight into how they come about, and assigns determinate probabilities to them. These probabilities are connected to the rest of our theory by laws that permit both prediction and (where means exist) control: if undisturbed, nucleus *a* will have probability *p* to alpha-decay (so we should expect *a*’s decay with *epistemic* probability *p*); and if we wish to alter *p*, our theory tells us how *a* must be disturbed.

It has been objected to the view of probability taken in this paper that unless probability attributions are interpreted as predictions about how relative frequencies will *actually* come out in the long run, probabilistic laws lack empirical content. Thus, if the relative frequency of decayed atoms in a large sample of some radioelement were, over a great length of time, to diverge significantly from the probability theoretically attributed to decay, that attribution would not be “borne out,” i.e., would be falsified. Otherwise, it is argued, probabilistic laws are compatible with all frequencies, and empirically vacuous.

But it is impossible for a world to “bear out” all of its probability-attributing laws in this sense. For these laws imply, among other things, that it is extremely unlikely that *all* actual long-run sequences will show a relative frequency near to the single-case probability. Therefore, the demand that all long-run decay rates nearly match all corresponding decay constants comes to a demand that nothing improbable show up in the long run, which is itself an improbability showing up in the long run. Intended to clear things up on the epistemological front, this proposal cannot even get out of its own way.

By splitting apart probabilistic explanation and induction, the *D-N-P* model has lost the point of probabilistic explanation. Behind this objection lies the view that probabilistic (or statistical) explanation is an activity fundamentally unlike *D-N* explanation. A probabilistic explanation is seen as a piece of detective work. Unable to give a causal demonstration of the explanandum from evidence thus far assembled, we develop hypotheses, which are judged by how probable they are on the evidence, and whether they make the explanandum sufficiently probable. In the end, we put forward the most convincing inductive argument yet found—the one making the explanandum most antecedently probable, given what else we know about events leading up to it.

This view of probabilistic explanation confuses epistemic with objective probability, and induction with explanation. Perhaps responsible for this confusion is the similarity of the tasks of explaining a phenomenon, gathering support for such an explanation, and gathering before- or after-the-fact evidence for a phenomenon’s occurrence. This confusion is abetted by misleading ways of talking about “strong” or “good” explanations. We should distinguish the following. (i) *A strong (good) explanation* is one that has great theoretical power, regardless of how well-confirmed it is or how probable it holds the explanandum to be. (ii) *A strong (good) candidate for explanation* is a proffered explanation with well-confirmed *premises*, regardless

of how probable it holds the explanandum to be and irrespective of how theoretically powerful it happens to be. (iii) *A strong (good) reason for believing that the explanandum fact will obtain* is furnished by before-the-fact evidence that leads, via one's theory, to an expectation of the explanandum with high epistemic probability. (iv) *A strong (good) reason for believing that the explanandum fact obtained* is given by any evidence that lends high epistemic probability to the proposition that the explanandum fact is a fact. Strong after-the-fact evidence, even for very improbable events, may be easy to come by. Reasons of types (iii) and (iv) need have nothing to do with explanation, and may be based on symptoms (Will it rain today?—Harry's rheumatism is acting up) or even less causally relevant information (Was Sue upset?—Her brother is certain she would have been).

Although the link between probabilistic explanation and induction is looser on the *D-N-P* model than on the *I-S* model, this is no fault: on Hempel's account it was entirely too close. Measuring the strength or "acceptability" of an explanation by the magnitude of the probability it confers on the explanandum blurs the distinctions just made. Keeping (i)–(iv) distinct, the *D-N-P* model enables us to state quite simply the object of induction in explanation: given a particular fact, to find, and gather evidence for, an explanans that subsumes it; given a generalization, to find, and gather evidence for, a higher-level explanans that subsumes it; in all cases, then, to discover and establish a true and relevant explanans. The issue of showing the explanandum to have high (relative or absolute) probability is a red herring, distracting attention from the real issue: the truth or falsity, and applicability, of the laws and facts adduced in explanatory accounts.

REFERENCES

- [1] Evans, R. D. *The Atomic Nucleus*. New York: McGraw Hill, 1965.
- [2] Giere, R. N. "Objective Single-Case Probabilities and the Foundations of Statistics." In *Logic, Methodology and Philosophy of Science, Vol. IV*. Edited by P. Suppes, et al. Amsterdam: North-Holland, 1973.
- [3] Hempel, C. G. "Aspects of Scientific Explanation." In *Aspects of Scientific Explanation and Other Essays*. New York: Free Press, 1965.
- [4] Hempel, C. G. "Maximal Specificity and Lawlikeness in Probabilistic Explanation." *Philosophy of Science* 35 (1968): 116–133.
- [5] Jeffrey, R. C. "Statistical Explanation vs. Statistical Inference." In *Essays in Honor of C. G. Hempel*. Edited by N. Rescher, et al. Dordrecht: D. Reidel, 1970.
- [6] Salmon, W. C. "Statistical Explanation." In *Statistical Explanation and Statistical Relevance*. Edited by W. Salmon. Pittsburgh: University of Pittsburgh Press, 1971.