

PETER RAILTON

PROBABILITY, EXPLANATION, AND
INFORMATION¹

INTRODUCTION

Elsewhere I have argued that probabilistic explanation, properly so called, is the explanation of things that happen by chance: the outcomes of irreducibly probabilistic processes.² Probabilistic explanation proceeds, I claimed, by producing a law-based demonstration that the explanandum phenomenon had a particular probability of obtaining and noting that, by chance, it did obtain. This account will be presented in more detail below, but the bulk of this paper will be given over to a discussion of two large problems confronting such a view: (1) many seemingly acceptable explanations of chance phenomena do not make explicit use of laws; and (2) many seemingly acceptable explanations that at least purport to be probabilistic concern phenomena known or assumed to be deterministic. Problem (1) is but an instance of a very general difficulty facing law-based accounts of explanation, and the approach to solving it suggested below may readily be extended to non-probabilistic explanation. Problem (2) raises a difficulty peculiar to probabilistic explanation, but I hope to show that an appropriate solution to problem (1) offers the key to a plausible solution to problem (2).

Discussions of both problems run the risk of degenerating into merely verbal quibbles. However, I believe that the issues concerning the application of the expressions 'explanation' and 'probabilistic explanation' which will be discussed here are genuine, and that we may gain understanding of explanatory practice by seeing how one might resolve problems (1) and (2). To revive an old way of putting it, there is a worthwhile analogy between 'explanation' and 'proof'.³ Various things have come to be called 'proofs'—deductive arguments, experiments, testimony, and testimonials, to name a few—but when well-informed we do not allow these various things to play the same roles in our lives. When we do logic, science, jury duty, or shopping, we attend to the relevant differences among so-called

“proofs”. Similarly, various things have come to be called ‘explanations’ – deductive arguments, inductive arguments, diagnoses, explications of meaning, excuses, and apologies, to name a few – but when well-informed we do not allow these various things to play the same roles in our lives. When we do science, engineering, politics, and auto repair we attend to the relevant differences among so-called “explanations”. ‘Probability’, too, is applied to various things – degrees of belief, degrees of confirmation, relative frequencies, chancy dispositions, and the outcome of combinatorics, among others. These differences, too, must be attended to when we look for guidance in our lives. I hope it will be clear in what follows that what one says about explanation and probabilistic explanation, and thus what one says about ‘explanation’ and ‘probabilistic explanation’, may make a difference that is a difference to actual practice.

I. THE DEDUCTIVE-NOMOLOGICAL-PROBABILISTIC ACCOUNT OF PROBABILISTIC EXPLANATION

The essence of probabilistic explanation, according to the deductive-nomological-probabilistic (D-N-P) account, is a deductive demonstration to the effect that an empirical theory assigned a particular probability to an explanandum phenomenon. If the relevant parts of the theory are true, the factual premises involved accurate, and the deduction valid, then the D-N-P explanation is true.⁴ If the premises are only partly true, or only approximately true, then the explanation may be only partly or approximately true. (Whatever partial and approximate truth are, this is not the place to analyze them.)

I take it that quantum mechanics, under the dominant interpretation, gives us reason to think that there are irreducibly probabilistic processes in nature, and that they are governed by probabilistic and non-probabilistic laws. Thus, the occurrence of a given alpha-decay, or of particular rates of alpha-decay, is widely held to be the result of a physically random process: spontaneous nuclear disintegration. If the dominant interpretation of quantum mechanics is right, there are no “hidden variables” characterizing unknown initial conditions of nuclei that suffice, in conjunction with deterministic laws, to account for the occurrences of alpha-decay: two radio-nuclei may be in the same physical state at a time t_0 , and may be subjected

to the same environment during the time interval from t_0 to $t_0 + \tau$ and yet one may decay during τ and the other not. Although this process is physically random in the sense that there is no "sufficient reason" for the one atom to decay but not the other, alpha-decay does obey physical laws: probabilistic laws such as those concerning barrier penetration and decay rate; deterministic laws such as those concerning the conservation of mass-energy and charge. Using certain simplifying assumptions, and substituting the appropriate values for decay energy, atomic number, and atomic weight in the Schrödinger wave-equation, it is possible to derive the decay constants of radioelements from quantum mechanics. Moreover, it has become conventional wisdom that such decay constants are not mere statistical averages drawn from large samples, but rather physically-irreducible, single-case probabilities—probabilities "per unit time for one nucleus".⁵ The psi-function of the wave-equation, then, is viewed as giving a complete description of the state of individual systems, and this state-description may at best determine only a probability distribution over future states.

For an example of a D-N-P explanation of a particular chance fact, let us suppose that a Uranium 238 nucleus u has alpha-decayed during a time interval of length τ which began at t_0 . Current physical theory assigns a decay constant, λ_{238} , to U^{238} nuclei, which allows us to state the probability that an individual U^{238} nucleus will decay during an interval of length t in the absence of external radiation: $(1 - \exp(-\lambda_{238} \cdot t))$. If we assume that u was not excited by external radiation at t_0 or after, we may make the following deductive-nomological inference:

- (1a) All U^{238} nuclei have probability $(1 - \exp(-\lambda_{238} \cdot \tau))$ to emit an alpha-particle during any interval of length τ , unless subjected to environmental radiation.
- (1b) u was a U^{238} nucleus at t_0 and was not subject to environmental radiation during the interval $t_0 - (t_0 + \tau)$.
-
- (1c) u had probability $(1 - \exp(-\lambda_{238} \cdot \tau))$ to alpha-decay during $t_0 - (t_0 + \tau)$.

Inference (1) yields a D-N explanation of the fact that u had a particular probability to alpha-decay at the time in question.

According to the D-N-P account, if (1) were prefaced with a

theoretical derivation of the law premise (1a) from underlying quantum physics, and were followed by a parenthetic addendum to the effect that u did indeed alpha-decay during the interval t_0 — $(t_0 + \tau)$, the result would be a full probabilistic explanation of u 's alpha-decay. Schematically, a D-N-P explanation of the fact that e is G at some time t_0 has the form (2) as follows:

- (2a) A theoretical derivation of a probabilistic law of the form
 (2b) $\forall x \forall t [F_{x,t} \rightarrow \text{Probability } (G)_{x,t} = r]$
 (2c) F_{e,t_0}
-
- (2d) $\text{Probability } (G)_{e,t_0} = r$
 (2e) (G_{e,t_0}) .⁶

Explanation (2) is a *purported* probabilistic explanation in D-N-P form. If its premises were true and its logic impeccable, it would be a *true* one as well.

It is obvious from the universal form of law (2a), and from the use of universal instantiation and *modus ponens* in deriving the conclusion (2d), that the probability involved is a single-case probability. The parenthetic addendum (2e) is the explanandum, and yet it in no way follows from the other premises. Is this paradoxical? No: it had better not be derivable in this way, for we are supposing that G_{e,t_0} occurred by chance, and thus that its occurrence cannot be derived from initial conditions by any empirical law. Presumably, we knew that G_{e,t_0} before any explanation was offered, and so step (2e) brings no news. In that sense, it is dispensable. Indeed, it is dispensable in another sense as well: those holding the view that explanations are always *arguments*—and should contain nothing inessential or incidental to such arguments—will want to see the addendum gone from the D-N-P schema. I have no real quarrel with either reason for dropping it, except that it will be of use in stringing together D-N-P explanations or amalgamating them with larger explanations; but more on this later. Moreover, I would like to urge that we move away from the view that explanations are purely arguments and pursue instead the idea that explanations are *accounts*—accounts in which arguments play a central role, but do not tell the whole story. Again, we will return to this below. For now I leave the parenthetic addendum in place, a token of my disaffection with the “pure argument” view.

It is more important to observe that the D-N-P account makes nothing of the *degree* of probability of the explanandum according to the explanans: in schema (2), r may take on any value in the unit interval without prejudice to explanation, as long, of course, as the probability-attribution is accurate. Thus, in contrast to Hempel's inductive-statistical model of probabilistic explanation, the D-N-P account does not require *high* probability either to establish the inductive relevance of the explanans to the explanandum (since deductive relevance takes its place), or to establish the nomic expectability of the explanandum (since nomic expectability "with practical certainty" is replaced by nomic expectability with whatever probability the explanandum had).⁷ This has the considerable advantage that we are no longer forced to say that less-than-highly-probable facts are inexplicable in principle. Some things that happen are vastly improbable, such as the decay at a given moment of an atom of U^{238} (mean life: 6.5×10^9 years), or are merely probable, such as the creation of a levorotatory form when an amino acid is synthesized (probability under standard conditions: .5); and when merely probable or vastly improbable things do occur, the appropriate explanation is that they had a particular probability of occurring, which chanced to be realized. Probabilistic explanation thus is not a second-string substitute for deterministic explanation, showing that the explanandum phenomenon *almost* had to occur. On the contrary, probabilistic explanation is a form of explanation in its own right, charged with the distinctive task of dealing with phenomena that came about by chance. As such, there is no special explanatory virtue in probabilistic explanations that show their explanandum to have been highly probable, unless the explananda in question *were* highly probable. If they weren't, then any explanations purporting to establish their high probability would not be explanatorily virtuous, but explanatorily inaccurate. Probabilistic explanations conferring high probability may have other sorts of virtues, however: other things equal, they receive greater inductive support from the evidence that the explanandum phenomenon occurred; other things equal, they may support more definite policies for the future; and so on. But we should not confuse all of the virtues of explanations with explanatory virtues. Without being a better explanation, one explanation may be more easily confirmed, more fully confirmed, more readily applied, or more easily translated into modern Greek than another.

A satisfactory D-N-P explanation need not even pick out factors that *raise* the probability of the explanandum over what it would have been in their absence.⁸ Consider the following example. Suppose that we are applying a herbicide to a patch of healthy milkweed, and suppose that a dose of this herbicide alters the biochemical state of milkweeds from a normal, healthy state *S*, in which plants have probability .9999 of surviving 24 hours, to a state *S'*, in which there is but probability .05 of lasting that long. When we return to the milkweed patch 24 hours after spraying and find to our consternation that a particular plant which received a full dose of herbicide is still standing, how are we to explain this? Presumably, we should point out that the plant was in state *S'* after the spraying, that in this state it had probability .05 of surviving 24 hours, and that, by chance, it did. The spraying, then, is part of the explanation of survival even though it *lowered* the probability of survival for this plant from what it would otherwise have been. We may not wish to speak of the spraying as a *probabilistic cause* of survival, since we may want to reserve the expression 'probabilistic cause' for factors that *do* raise the probability of an event in the circumstances. Thus, for those plants failing to survive, we could speak of the spraying as a probabilistic cause of their deaths, while for the plant that survived, it would be strained at best to speak of the spraying as a probabilistic cause of its survival. This suggests that probabilistic explanation is not a mere subspecies of causal explanation.⁹

It might be objected that if probabilistic explanations do not establish high probability for their explananda or do not point to factors that raise the probability of their explananda, then how do they show why one outcome – alpha-decay, survival, or whatever – was realized rather than its opposite? The short answer is that they do not show this and should not try. If there were a reason why one probabilistic outcome of a chance process was realized rather than another, we would not be dealing with a chance process.

Finally, because the D-N-P account is deductive in form and based upon lawful, physical, single-case probabilities, it can be shown not to suffer from the problems of explanatory ambiguity and epistemic relativity that have bedeviled the Hempelian inductive-statistical account.¹⁰ If the very same fact were to be assigned two different lawful, physical, single-case probabilities by our theory, then contradiction, not ambiguity, would have to be dealt with, and we should

proceed accordingly. As for epistemic relativity, incomplete knowledge may make something *look like* a genuinely probabilistic law or a genuine probabilistic explanation, but it cannot make things *be* other than they are. This is strictly parallel with non-probabilistic explanation: in both cases there are problems about distinguishing laws from mere factual regularities; in both cases there are problems about distinguishing the true from the false. This is not to say that there are no special epistemological problems in the probabilistic case, but only to say that these problems *are* epistemological and do not make successful explanation independent of either lawfulness or truth.

With this brief sketch of the D-N-P account in hand, let us turn to two important problems it faces.

II. PROBABILISTIC EXPLANATIONS WITHOUT LAWS

A difficulty that arises at once for the D-N-P account is that many of the probabilistic explanations offered and accepted in scientific discourse fall far short of providing all of the elements of schema (2). We have encountered an example already in the explanation given of a particular milkweed plant's survival:

- (3) "the plant was in state *S'* after the spraying, ... in this state it had probability .05 of surviving 24 hours, and ... , by chance, it did".

Explanation (3) contains no theoretical derivation, no explicit law, and yet seems highly acceptable. Other less-than-full-fledged probabilistic explanations encountered in scientific contexts are bare indeed: "Why did this muon decay? – Because it was unstable"; "Why is the Geiger counter clicking as we approach that rock? – Because it contains uranium"; "Why is this one lobster blue? – It's a random mutation, very rare"; and so on. Proffered probabilistic explanations may fail to include all of the elements of (2) for a variety of reasons: in certain contexts, a more elaborate explanation may be out of place – the audience may be too well-versed, not well-versed enough, not interested enough, or short on time; a more elaborate explanation may not be available even if it were appropriate – the relevant laws and facts may not be known, or may be known only qualitatively; the person offering the explanation may simply not know enough; and so

on. Whatever the reason, does the D-N-P account force us to say that these less-than-full-fledged specimens are not explanations?

In many cases, that would be an intolerably strict position to take. But where should one draw the line between explanation and non-explanation? The answer lies in not drawing lines, at least at this point, and in recognizing instead a continuum of explanatoriness. The extreme ends of this continuum may be characterized as follows. At one end we find what I will call an *ideal D-N-P text*, comprising all that schema (2) involves. At the other end we find statements completely devoid of what I will call "explanatory information".¹¹ What is explanatory information? Consider an ideal D-N-P text for the explanation of a fact *p*. Now consider any statement *S* that, were we ignorant of this text but conversant with the language and concepts employed in it and in *S*, would enable us to answer questions about this text in such a way as to eliminate some degree of uncertainty about what is contained in it. To the extent that *S* enables us to give accurate answers to such questions, i.e., to the extent that it enables us to reconstruct this text or otherwise illuminates the features of this text, we may say that *S provides explanatory information concerning why p*. It is hardly novel to speak of sentences conveying information about complete texts in this way: presumably we employ such a notion whenever we talk of a piece of writing or an utterance as a summary, paraphrase, gloss, condensation, or partial description of an actual text, such as a novel, a speech, or a news report. Unfortunately, I know of no satisfactory account of this familiar and highly general notion, and so cannot appeal to one here. Nor can I begin to provide an account of my own making within *this* text. To make matters worse, I cannot even pretend to be using the well-defined notion of information employed in Wiener-Shannon information theory. However, before I say anything more about the sort of analysis that is needed for the concept of information employed here, let us look at a few examples to see of what use that concept might be.

To be told that a Geiger counter is clicking because it is near a uranium-bearing rock is to be given explanatory information—assuming that the claim is true as far as it goes—since it would be part of the relevant ideal D-N-P text for the clicking that the rock in question contains uranium. Of course, the ideal D-N-P text would go on to tell us about the radioactivity of uranium, about the process of alpha-decay, about the mechanics of Geiger counters, and so on.

Obviously, the amount of explanatory information provided by the brief and informal explanation given above is much less than the amount needed to reconstruct a sizeable chunk of the entire ideal D-N-P text, sufficing only to illuminate one central feature of the whole text. But to illuminate one central feature is to convey a not insignificant amount of information, and thus to be off the zero point of our continuum. Somewhat nearer that point, but still above it, is the explanation "Why did this muon decay?—Because it was unstable". This explanation accurately points to a dispositional property of the muon responsible for the decay, distinguishing this decay from disintegration due to external forces such as bombardment by other particles; and as we know from the shortcomings of extensionalist accounts of dispositional properties, it is one thing to say that the muon decayed, and something else—something more—to say that it decayed owing to a disposition to do so. The disposition in this case is probabilistic, and so the relevant ideal explanatory text would be D-N-P in form. The brief statement about instability thus answers certain questions about the form as well as the content of the relevant ideal text. The explanation of a lobster's blueness in terms of a "random mutation, very rare" also tells us that the relevant explanatory text is D-N-P in form, and further indicates the particular probabilistic process involved—genetic mutation—and provides a qualitative characterization of the probability of blueness. Thus it falls higher on the explanatory continuum than the Geiger counter clicking or muon decay explanations. A bit higher on the continuum is what (3) says about the explanation of a particular milkweed plant's survival. (3), too, tells us that the relevant ideal text is D-N-P in form, and mentions an important factor in the process culminating in the explanandum—the spraying—but it further supplies a precise, quantitative evaluation of the probability involved. Much higher on the continuum are the explanatory remarks made earlier about the alpha-decay of U^{238} nucleus u , for they contained a probabilistic law explicitly and roughed out a theoretical derivation of this law. Higher still are the well-developed explanatory accounts one finds in physics or chemistry textbooks and monographs. These are not usually explanations of particular facts, but our characterization of explanatory information is indifferent as to whether p is a particular or general fact. Moreover, outside the context of standard teaching and reference works, scientists are often concerned with

particular-fact explanation – e.g., in the course of experimentation – and, as the D-N-P account recognizes, elaborate theoretical derivations have a natural place in fully developed explanations of particular facts.

Perhaps this is the point at which to say something more about what the theoretical derivations figuring in schema (2) might look like. The place to look for guidance is plainly scientific explanatory practice itself. If one inspects the best-developed explanations in physics or chemistry textbooks and monographs, one will observe that these accounts typically include not only derivations of lower-level laws and generalizations from higher-level theory and facts, but also attempts to *elucidate the mechanisms* at work. Thus an account of alpha-decay ordinarily does more than solve the wave-equation for given radionuclei and their alpha-particles; it also provides a model of the nucleus as a potential well, shows how alpha-decay is an example of the general phenomenon of potential-barrier penetration (“tunnelling”), discusses decay products and sequences, and so on. Some simplifying assumptions are invariably made, along with an expression of hope that as we learn more about the nucleus and the forces involved we will be able to give a more realistic physical model. It seems to me implausible to follow the old empiricist line and treat all these remarks on mechanisms, models, and so on as mere *marginalia*, incidental to the “real explanation”, the law-based inference to the explanandum. I do not have anything very definite to say about what would count as “elucidating the mechanisms at work” – probabilistic or otherwise – but it seems clear enough that an account of scientific explanation seeking fidelity to scientific explanatory practice should recognize that part of scientific ideals of explanation and understanding is a description of the mechanisms at work, where this includes, but is not merely, an invocation of the relevant laws. Theories broadly conceived, complete with fundamental notions about how nature works – corpuscularianism, action-at-a-distance theory, ether theory, atomic theory, elementary particle theory, the hoped-for unified field theory, etc. – not laws alone, are the touchstone in explanation. Of course, there *are* marginal comments to be found accompanying the explanations offered in scientific textbooks and monographs: one reads of helpful ways of visualizing, conceptualizing, applying mathematical devices, etc., which the reader is warned not to take too seriously. There is no hard and fast line between simplified

idealizations of actual physical processes – such as the model of a gas as a collection of perfectly elastic spheres or the “liquid-drop” model of the atomic nucleus – and more genuine marginalia or heuristics – such as the treatment of electron orbits as harmonic vibrations of strings, the analogy between electrostatics and hydrostatics, the use of topological images in describing systems with multiple equilibria, and so on. Moreover, one theory’s attempt at a realistic model may serve as no more than a first-order approximation for later developments of that theory – e.g., atomic theory in this century. Genuine marginalia and heuristics are usually flagged as such, but in a developing theory there may be indefiniteness about whether a given model should be taken in a realistic and explanatory way or as an unrealistic “picture” with instrumental value only. No analysis of explanation should be more definite on such questions than the best available theories themselves; so while we should ask of scientists whether a particular model ought to be part of ideal explanatory texts, we should be prepared to accept indefiniteness and even disagreement in the answers we receive.

At a particular time we may have nothing like the knowledge necessary to produce an ideal D-N-P text for a given chance phenomenon. Not knowing fully what the relevant ideal text looks like, we evidently will be unable in many cases to say how much explanatory information a given proffered explanation carries. But again, it is not the job of an analysis of explanation to settle questions beyond the reach of existing empirical science. Instead, it is appropriate and adequate that an analysis of explanation define what *would be* explanatory information in such cases. On the analysis given here, a proffered explanation supplies explanatory information (whether we recognize it as such or not) to the extent that it does in fact (whether we know it or not) correctly answer questions about the relevant ideal text. Whether in a given context we *regard* a proffered explanation as embodying explanatory information, in light of the interpretation we impose on it and our epistemic condition generally, is a matter for the pragmatics of explanation.

The general form of the ideal D-N-P text is meant to represent the ideal striven for in actual explanatory practice, i.e., it comprises the things that a research program seeks to discover in developing the capacity to produce better explanations of chance phenomena. Thus, the ideal D-N-P text reflects not only an ideal of explanation, but of

scientific understanding: we may say that we understand why a given chance phenomenon occurred to the extent that we are able, at least in principle, to produce the relevant ideal D-N-P text or texts. However, the concept of scientific understanding – and its links to ideals of explanation and to fundamental conceptions of nature – deserve more serious treatment than it can be given here. Needless to say, even if we did possess the ability to fill out arbitrarily extensive bits of ideal explanatory texts, and in this sense thoroughly understood the phenomena in question, we would not always find it appropriate to provide even a moderate portion of the relevant ideal texts in response to particular why-questions. On the contrary, we would tailor the explanatory information provided in a given context to the needs of that context; if we had the capacity to supply arbitrarily large amounts of explanatory information, there would be no need to flaunt it.

As mentioned earlier, there is no ready-made analysis of information capable of doing the work here asked of the concept of information. The sort of information discussed in *information theory* is *syntactic*, a measure of the statistical rarity of signals from an observed source, and the sort discussed here is *semantic*, it is information *about* something (an ideal text).¹² However, there is an important analogy that makes the same term applicable: in both cases the amount of information carried by a “message” is proportional to the degree to which it reduces uncertainty; thus a statistically common signal from a stable source embodies less information than a statistically rare one, and a fuller proffered explanation leaves less to be known about the relevant ideal text than a slighter one. Hence, information is a kind of *selection power* over possibilities. But here the analogy ends. Semantic equivalents, such as ‘son’ and ‘male offspring’, *as signals* may require different amounts of information to be transmitted; and there is no guarantee whatsoever that a message embodying huge amounts of syntactic information (and so requiring a great deal of channel capacity to transmit) will embody much if any semantic information – witness certain television test patterns (and perhaps certain programs as well). It is rather elusive of me to speak of explanatory information without having an adequate theory of semantic information, but none of us has an adequate theory of semantic information and the notion still receives very heavy use. I trust that as I join the multitude it is intuitively (deceptively?) clear what this notion involves.

Perhaps a few things might be said by way of clarification. If a self-consistent sentence S implies a sentence S^* , then S contains at least as much semantic information as S^* . Thus, for example, ' p ' is at least as informative as ' $p \vee q$ ', since ' p ' reduces the range of possibilities at least as much as ' $p \vee q$ '; and ' $p \& q$ ' is at least as informative as ' p ' or ' q ' alone, since the conjunction reduces the range of possibilities at least as much as either individual conjunct. If S is a tautology, then S contains no semantic information (at least, if we take logic for granted), since S eliminates no (logical) possibilities. If S contains semantic information $I(S)$, and S^* contains semantic information $I(S^*)$, and if the conjunction of S and S^* is self-consistent, then the amount of semantic information contained in the conjunct will equal $I(S)$ plus $I(S^*)$ minus any overlap. And so on. In these properties of semantic information it may be seen that, as with syntactic information, this quantity bears something of an inverse relation to probability. Propositions with logical probability one, tautologies, convey zero semantic information; conjunction raises semantic information, but for independent events it lowers probability; disjunction lowers semantic information, but for independent events it raises probability; and so on. But a thicket of problems confronts any attempt to press an analysis of semantic information, or its formal parallels with syntactic information, much further. First, if we are going to speak of semantic information as reducing uncertainty, we must have a probability measure of the appropriate kind. The Carnap-Bar-Hillel approach to semantic information uses Carnapian logical probability, but this brings with it grave problems.¹³ Second, to talk of the semantic content of a sentence is to assume some *semantic interpretation* of it, and this raises a number of issues: What sort of semantics should we use? Must semantic interpretation be a function of context? What constraints must be met in order for a semantic interpretation to exist? And so on. For our purposes I propose to ignore the general questions about semantic interpretation and speak only of uncertainty as a measure of the extent to which, in what I will call a "standard context", one is unable to reconstruct particular or general features of the relevant ideal explanatory text. Clearly, any general analysis of semantic information would have to provide a measure of the extent to which a given sentence allows us, in a *variety* of contexts, to "recover" or "reconstruct" an aspect of *the world* itself. It is less problematic, though by no means easy, to treat of semantic information in a particular *kind* of context and with

respect to a definite *text*. The sort of *standard context* I have in mind is the following: one understands all of the concepts and terms used in the relevant ideal text and in the proffered explanation, but is ignorant of the overall structure and the details of that ideal text. Hereinafter, when I speak of a proffered explanation conveying explanatory information, I will mean that it would do so in the appropriate standard context.¹⁴

The account of explanatory information just outlined is liable to trivialization. For example, suppose that an alpha-decay occurs, and suppose that someone proffers the following "explanation": "the relevant ideal text contains more than 10^2 words in English". For all that has been said here, this remark, if true (as it doubtless is), would count as conveying explanatory information and hence as some sort of an explanation. The first line of defense against this kind of degenerate case would be to say that the information must be about the *content* and not merely the *form* of the relevant ideal text. But this is not satisfactory, for it certainly is explanatory to be told, of the alpha-decay, that the relevant ideal text is probabilistic rather than deterministic in form. In order to avoid further complication, I will simply tolerate this kind of degenerate case, relegating it to the very low end of the continuum of explanatoriness. And rightfully enough: it certainly does not afford much illumination of the relevant ideal explanatory text.

In sum, proffered explanations of chance phenomena that fail to live up to the standards of the D-N-P schema may still count as explanations, even as quite good explanations, in virtue of finding another way of communicating explanatory information. In context, such explanations may be more than adequate.

What has been said here about the potential explanatoriness of statements lacking strict D-N-P form can be extended in obvious ways to non-probabilistic explanations and may be used to replace the justly criticized idea that the D-N schema provides necessary conditions for successful explanations. I would argue that the D-N schema instead provides the skeletal form for ideal explanatory texts of non-probabilistic phenomena, where these ideal texts in turn afford a yardstick against which to measure the explanatoriness of proffered explanations in precisely the same way that ideal D-N-P texts afford a yardstick for proffered explanations of chance phenomena. Thus, proffered explanations of non-probabilistic phenomena may take

various forms and yet still be successful in virtue of communicating information about the relevant ideal text. For example, an ideal text for the explanation of the outcome of a causal process would look something like this: an inter-connected series of law-based accounts of all the nodes and links in the causal network culminating in the explanandum, complete with a fully detailed description of the causal mechanisms involved and theoretical derivations of all the covering laws involved. This full-blown causal account would extend, via various relations of reduction and supervenience, to all levels of analysis, i.e., the ideal text would be closed under relations of causal dependence, reduction, and supervenience. It would be the whole story concerning why the explanandum occurred, relative to a correct theory of the lawful dependencies of the world. Such an ideal causal D-N text would be infinite if time were without beginning or infinitely divisible, and plainly there is no question of ever setting such an ideal text down on paper. (Indeed, if time is continuous, an ideal causal text might have to be non-denumerably infinite – and thus “ideal” in a very strong sense.) But it is clear that a whole range of less-than-ideal proffered explanations could more or less successfully convey information about such an ideal text and so be more or less successful explanations, even if not in D-N form.

It is preposterous to suggest that any such ideal could exist for scientific explanation and understanding? Has anyone ever attempted or even wanted to construct an ideal causal or probabilistic text? It is not preposterous if we recognize that the actual ideal is not to *produce* such texts, but to have the ability (in principle) to produce arbitrary parts of them. It is thus irrelevant whether individual scientists ever set out to fill in ideal texts as wholes, since within the division of labor among scientists it is possible to find someone (or, more precisely, some group) interested in developing the ability to fill in virtually any particular aspect of ideal texts – macro or micro, fundamental or “phenomenological”, stretching over experimental or historical or geological or cosmological time. A chemist may be uninterested in how the reagents he handles came into being; a cosmologist may be interested in just that; a geologist may be interested in how those substances came to be distributed over the surface of the earth; an evolutionary biologist may be interested in how chemists (and the rest of us) came into being; an anthropologist or historian may be interested in how man and material came into

contact with one another. To the extent that there are links and nodes, at whatever level of analysis, which we could not even in principle fill in, we may say that we do not completely understand the phenomenon under study.

The full explanatory history of a given phenomenon may have both causal and probabilistic links; for example, the click of a Geiger counter placed near a rock is due to geological causes as well as spontaneous alpha-decay. There is no problem in saying that an ideal explanatory text for a given fact contains both D-N and D-N-P subparts, so long as we do not try to treat the very same links as both probabilistic and non-probabilistic.

As we noticed earlier, in cases of genuine probabilistic explanation there are certain why-questions that simply do not have answers – questions as to why one probability rather than another was realized in a given case. We now are in a position to say what this comes to: such why-questions are requests for information that simply is not available – no part of even the ideal explanatory text contains a sufficient reason why one probability was realized rather than another. That is, this request for further explanation is refused, not because we do not know enough, but because there is simply nothing more to be known.¹⁵ Nor is this the only sort of case in which a why-question may be refused. In any theory, probabilistic or otherwise, there will be certain matters of ultimate brute fact or of absolutely fundamental law. Purported ideal explanatory texts based upon a theory may simply lack the material needed to provide answers to certain requests for further explanatory information. For example, current physics offers us no reason why there is no negative gravity, despite the fact that there is both negative and positive charge (perhaps this is an initial condition of our universe), nor does it offer a reason why mass-energy is conserved rather than not (this seems to be a fundamental law). At most basic level, it may be difficult to distinguish fundamental laws from initial conditions of the entire universe, but that is not the point at issue here. The point is that a theory may legitimately spurn certain requests for further explanation in both probabilistic and non-probabilistic cases, and we can now say what such spurning consists in: the *absence* of certain things from purported ideal explanatory texts based upon that theory.

Where the orthodox covering-law account of explanation propounded by Hempel and others was right has been in claiming that explanatory practice in the sciences is in a central way *law-seeking* or

nomothetic. Where it went wrong was in interpreting this fact as grounds for saying that any successful explanation must succeed either in virtue of explicitly invoking covering laws or by implicitly asserting the existence of such laws. It is difficult to dispute the claim that scientific explanatory practice—whether engaged in causal, probabilistic, reductive, or functional explanation—*aims* ultimately (though not exclusively) at uncovering laws. This aim is reflected in the account offered here in the structure of ideal explanatory texts: their backbone is a series of law-based deductions. But it is equally difficult to dispute the claim that many proffered explanations succeed in doing some genuine explaining *without* either using laws explicitly or (somehow) tacitly asserting their existence. This fact is reflected here in the analysis offered of explanatoriness, which is treated as a matter of providing accurate information about the relevant ideal explanatory text, where this information may concern features of that text other than laws. For definiteness, let us call the analysis of explanation offered here *the nomothetic account of scientific explanation*.

III. THE USE OF STATISTICS AND PROBABILITIES IN THE EXPLANATION OF NON-PROBABILISTIC PHENOMENA

Now that the nomothetic account of probabilistic and causal explanation has been introduced, we are in a position to deal with the second of the two problems posed at the outset. Thus far it has been assumed that the proffered probabilistic explanations we have considered are purported explanations of indeterministic phenomena, so that the relevant ideal texts (or the relevant portions of more comprehensive ideal texts) are D-N-P in form. But what if a seemingly probabilistic explanation is offered of a phenomenon brought about in a deterministic way? Must we dismiss all such explanations as worthless? Would not this make many important statistical explanations in science—e.g., in classical statistical mechanics—worthless? But surely we cannot say *that*.

Fortunately, there is no need to make such a sweeping dismissal of statistical explanations of non-random phenomena. For while statistical explanations of phenomena not due to chance cannot be explanatory in virtue of providing information about a relevant ideal

D-N-P text, since there is *no* such text, they may be explanatory in virtue of providing information about a relevant ideal D-N text, since there *is* just such a thing. To see this, let us look briefly and informally at two important cases of statistical explanation of non-random phenomena: classical statistical mechanics and “explanation by correlation”.

The foundations of statistical mechanics evoke a wealth of difficult issues that cannot be considered here. Moreover, classical statistical mechanics has been replaced by quantum statistical mechanics, so it is no longer accurate to say that thermodynamic phenomena (for example) do not involve real indeterminacy at all. To simplify matters, I will stick to the classical case, which will enable me to discuss a view that has been and continues to be widely held: if classical statistical mechanics *were* right and no indeterminacy were involved in thermodynamic phenomena, the statistical explanations of the classical theory would be genuinely explanatory. I have in mind familiar explanations of such facts as the tendency of closed systems to move toward and stay near equilibrium, the functional correlation of macroscopic thermodynamical variables for gases at equilibrium (e.g., the ideal gas law, $PV = kT$), etc., in terms of the most common micro-states in an ensemble created by permutations of molecules over a constrained phase space (Boltzmann) or in terms of average values over all possible micro-states in a constrained phase space (Gibbs). For any given gas, its particular state S at a time t will be determined solely by its molecular constitution, its initial condition, the deterministic laws of classical dynamics operating upon this initial condition, and the boundary conditions to which it has been subject. Therefore, the ideal explanatory text for its being in state S at time t will not be probabilistic, but will be a complete causal history of the time-evolution of that gas. Now if the ideal explanatory text is purely deterministic, how could one explain the fact that the gas is in state S at time t by reference to “most probable” macrostates or “phase averages”? After all, these “probabilities” and “averages” are the product of sheer combinatorics under certain limiting assumptions, with no guarantee that the gas in question met these limiting assumptions and with no regard to the fact that the gas had definite initial conditions, not a probability distribution over initial conditions. Briefly, the answer is that such appeals to combinatorics serve to illuminate a significant feature of the causal process underlying the behavior of classical thermodynamic systems, thereby serving to

partially illuminate the relevant ideal explanatory texts in particular cases. For example, various proofs in ergodic theory and related results show that if a gas is in an initial condition that obeys a relatively few constraints, it will, over infinite time, spend most of its time at or near equilibrium. This illuminates a *modal* feature of the causal processes involved and therefore a modal feature of the relevant ideal explanatory texts: this sort of causal process is such that its macroscopic outcomes are remarkably insensitive (in the limit) to wide variations in initial microstates. The stability of an outcome of a causal process in spite of significant variation in initial conditions can be informative about an ideal causal explanatory text in the same way that it is informative to learn, regarding a given causal explanation of the First World War, that a world war would have come about (according to this explanation) even if no bomb had exploded in Sarajevo. This sort of robustness or resilience of a process is important to grasp in coming to know explanations based upon it. Traditional worries about the applicability of ergodic proofs and related results to systems finite in time and degrees of freedom do not go away, on this view. But for those who have held that these worries can be dealt with, we now have a way of capturing the intuition that ergodic theory and its kin are somehow explanatory: they shed light on a modal feature of the casual processes underlying thermodynamic behavior, thus providing information about the relevant ideal causal texts. Since we almost never know the actual initial condition of a gas in any detail, this sort of modal information about the causal process involved in the time-evolution of classical gases is especially valuable.

Statistical generalizations about macro- and micro-states of thermodynamic systems—e.g., “closed systems are seldom found in states with measure zero (in the natural measure)” —also may play a role in explaining thermodynamic behavior, such as the prevalence of equilibrium. However, the role they play is not that of illuminating lawful features of the processes underlying such behavior, but rather that of providing information about the factual premises of the relevant ideal explanatory texts. Thus, it should be no mystery that statistical generalizations believed *not* to reflect underlying probabilistic laws may still be explanatory, once we recognize that they function in explanation not as ersatz laws but as summaries of information about initial conditions and boundary conditions.

It would be excessively fastidious to insist that classical statistical

mechanics offers no probabilistic explanations only if this were tantamount to insisting that it could offer *no* explanations, i.e., could provide no explanatory information. We have seen that classical-statistical-mechanical arguments *can* provide explanatory information concerning purely non-probabilistic processes. Thus classical statistical mechanics need hardly be dismissed as non-explanatory, although one must be careful in stating what it does and does not explain. Most crucially, it does not explain why the initial conditions should be such that we observe the macroscopic regularities in thermodynamic behavior that we do. This shows up on the present account in the fact that these initial conditions can be traced backwards through ideal causal texts only to still earlier conditions, not to laws. Thus the prevalence of equilibrium and other such features of macroscopic behavior must, on the classical theory, ultimately be attributed to brute fact and to the operation of deterministic laws on brute fact. But *given* a certain range of initial conditions, classical statistical mechanics does provide explanations of macroscopic behavior. The charge of excessive fastidiousness seems doubly unearned in light of the fact that classical statistical mechanics itself tells us that the relevant ideal explanatory texts are deterministic.

The second sort of statistical explanation that is frequently offered of phenomena known or believed not to be indeterministic falls under the general (often pejorative) heading "explanation by correlation". To this species belong the many efforts to use the direct results of multivariate analysis and other statistical techniques as explanations of the behavior of (often complex) causal systems. It is unnecessary in this connection to delve into the details of explanations of this kind, for the point to be made is highly general: the nomothetic account enables us to see how statements of statistical correlations may function not only as evidence for causal connections, but also as sources of information about aspects of ideal causal explanatory texts. Thus, a statement that there is a "statistically significant" correlation between exposure to high levels of cotton fiber in the air and incidence of brown lung disease may be put forward to explain why textile workers in cotton mills experience abnormally high rates of brown lung. Here the explanation may involve an irreducibly probabilistic process, but it need not. Let us suppose for the sake of illustration that it does not; still, the statement of correlation may convey information about the relevant causal ideal text for explaining the frequency of brown lung among

cotton mill workers, since it points to a substance, cotton fiber, the presence of which in the air breathed is an important element of this text. In this light, claims about the "strength" of a correlation, or about the amount of variation in one variable that can be "statistically explained" in terms of the variation of another, may be seen as ways of providing information about the extent and independence of the roles of various factors in relevant ideal explanatory texts.

It would be a naive view to equate the use of statistical correlations with probabilistic explanation, and we should not accuse scientists making heavy use of correlations – often because of limitations in the available evidence, limitations on possible experimentation, and the sheer complexity of the systems studied – of offering probabilistic explanations of phenomena they suspect are deterministic. "Explanation by correlation", of course, does not always provide genuinely explanatory information – correlations may be accidental, epiphenomenal, mediated by other variables, misleading about the direction of causation, or otherwise explanatorily defective. But even so, some such correlations may provide a limited amount of information about the relevant ideal explanatory texts. Thus, failure to detect an intervening variable may weaken an explanation greatly, but that explanation may still pick out factors with some causal role even if it misrepresents the directness and importance of their contribution. Clearly, "explanation by correlation" needs further and deeper treatment than it can be given here. But I hope that it is equally clear how the nomothetic account enables us to say that statements of statistical correlation may function as (partial) explanations in virtue of providing (partial) information about relevant ideal explanatory texts, even when these are not probabilistic.

Finally, while it is important to recognize that statistical generalizations of various sorts may provide explanatory information concerning non-probabilistic phenomena, it is important as well to recognize the difference it makes to explanatory practice whether or not a given phenomenon is viewed as being fundamentally probabilistic. For example, quantum mechanics, with its underlying probabilism, initially confronted a resistance from scientists quite unlike that meeting classical statistical mechanics, with its underlying determinism. Furthermore, there are differences in the research programs supported by these two theories: in the latter case we find sustained work in physics and mathematics to account for statistical

features of systems in terms of “hidden variables”, in the former case we find – after considerable debate – widespread acceptance of the non-existence of “hidden variables” and work in physics and mathematics to demonstrate the “completeness” of the quantum theory. These differences in the history of theory-acceptance and theory-development are the sort of thing that a satisfactory account of scientific explanation ought to illuminate, but an account of probabilistic explanation that fails to demarcate “statistical” explanations of deterministic processes (as in classical statistical mechanics) from probabilistic explanations of indeterministic processes (as in quantum mechanics) tends to obscure rather than clarify these differences. Unlike many of its competitors, the nomothetic account of probabilistic explanation – based as it is on the D-N-P model – makes precisely this demarcation, and in so doing helps us to understand these features of scientific explanatory practice. In this way, an account of explanation may itself do some explanatory work by accounting for a difference in practice in terms of a difference in theory. At first blush, one might think that whenever statistics or probabilities are involved in explanatory practice one is dealing with a form of probabilistic explanation. However, this illusion is quickly shed once one recognizes the variety of ways in which statistics and probabilities figure in explanatory activities. Perhaps the commonest use of statistics and probabilities in connection with explanation is *epistemic*: they are used in the process of assembling and assessing evidence for causal and non-causal explanations alike. Somewhat less common, but still important, are such uses as those discussed above: statistics and probabilities are used in providing explanatory information about causal and non-causal processes and their initial conditions. In *some* cases of the last sort we have genuine probabilistic explanation, specifically, in those cases where information is provided about a physically indeterministic process.

Some have protested that accepting the existence of fundamentally probabilistic phenomena amounts to a resignation of intellectual responsibility, to an abandonment of curiosity. Plainly it would be irresponsible to accept without reservation the hypothesis that a process is probabilistic on the basis of scanty examination and weak evidence, but it would be irresponsible to accept *any* hypothesis without reservation in such circumstances. And surely no one can argue that acceptance of the irreducible probabilism of quantum mechanics is based upon scanty examination or weak evidence.

Moreover, as we noted earlier, the explanations offered by a number of historically significant scientific theories must come to an end somewhere. Acceptance of a lawful probabilistic relationship as fundamental on the basis of overwhelming evidence is no greater bar to the growth of science than is acceptance of a lawful deterministic relationship as fundamental on the basis of equally good evidence. Certainly the history of quantum mechanics provides stronger support for this claim than any *a priori* argument possibly could: the deepening and widening of scientific knowledge under quantum theory gives the lie to those who prophesied that explanation by chance would be the enemy of enlarged inquiry and understanding.

The University of Michigan, Ann Arbor

NOTES

¹ I am very grateful to David Lewis, Lawrence Sklar, and Timothy McCarthy for helpful comments on some of the material contained in this paper. Some of the research for this paper was supported by a dissertation-writing fellowship from the Whiting Foundation at Princeton University and by a Junior Fellowship in the Society of Fellows at the University of Michigan, Ann Arbor; their support is acknowledged with thanks.

² In 'A Deductive-Nomological Model of Probabilistic Explanation', *Philosophy of Science* 45 (1978), 206–226.

³ This analogy has been most fully developed by C. G. Hempel. See his 'Aspects of Scientific Explanation', in *Aspects of Scientific Explanation and Other Essays* (New York: Free Press, 1965).

⁴ Here I follow Hempel's use of 'true' for explanations. See 'Aspects', p. 338.

⁵ R. D. Evans, *The Atomic Nucleus* (New York: McGraw-Hill, 1965).

⁶ The connective '→' is to be read in whatever way it should be read for laws of nature in general, e.g., as a "strong conditional". The parentheses enclosing the addendum (2e) are meant to indicate that it is not a logical consequence of the steps above it. More will be said about the nature of the theoretical derivation (2a) in what follows.

⁷ See Hempel, 'Aspects', section 3; and Railton, section 5.

⁸ Just how this sort of counterfactual is to be read is too large an issue to enter into in this context. Instead, in the example that follows I will assume what I take to be a natural reading of the relevant counterfactual.

⁹ I have borrowed this example of a plant surviving herbicide spraying from Nancy Cartwright (in conversation), who used it to make a different point. I am, of course, assuming that the effect of the herbicide is genuinely probabilistic – perhaps owing its indeterminism to the chance factors (such as electron location) that influence chemical bonding.

¹⁰ See, for example, Hempel, 'Aspects', section 3, and 'Maximal Specificity and Lawlikeness in Probabilistic Explanation', *Philosophy of Science* 35 (1968), pp. 116–133.

¹¹ I owe the expression 'explanatory information' to David Lewis, who, in the course of advising my dissertation on explanation, supplied me with this notion as a way of expressing the idea of illuminating ideal explanatory texts. I do not know whether he would agree with the way I am using this notion here.

¹² For a discussion of the differences between "information about" and Wiener-Shannon information, see Colin Cherry, *On Human Communication*, second edition (Cambridge, Massachusetts: MIT Press, 1966), chapter 6.

¹³ See Y. Bar-Hillel and R. Carnap, 'Semantic Information', *British Journal for the Philosophy of Science* 4 (1953), 147-157. It should be noted that a proffered explanation may communicate some explanatory information in virtue of the *organization* of its components as well as their content proper. Thus a tautology by itself lacks semantic information, but it may play a role in the structure of a proffered explanation that informatively reflects the structure of the relevant ideal text. There is no obvious way of incorporating *organization* into existing accounts of semantic information.

¹⁴ One particularly difficult problem for a theory of semantic information, especially one that focuses on information as selection power, will be to capture some form of the *de re/de dicto* distinction. For example, we may have two proffered explanations of a given phenomenon. One correctly but sketchily describes the relevant ideal explanatory text; the other is a disjunction of a very full description of the relevant ideal explanatory text and a feeble description of some quite irrelevant ideal explanatory text. Both explanations would then be true, and we can imagine that (by some intuitive measure) both manage overall to exclude an equally large number of irrelevant possibilities - the specificity of the first disjunct of the latter making up for the failure to eliminate the second disjunct. It may seem that the former explanation is more explanatory because it is more *about* the relevant ideal text than the latter, since the former allows us to identify at least some part of the relevant ideal text, while the latter does not. Obviously this is part of a general problem about *de re* and *de dicto*; depending upon one's approach to this problem, one may come up with different accounts of semantic information. I am grateful to Timothy McCarthy for drawing my attention to this issue.

¹⁵ B. C. Van Fraassen is among those who have argued that there is need for an analysis of the notion that certain requests for information may be refused.

REFERENCES

- Bar-Hillel, Y. and Carnap, R.: 1953, 'Semantic Information', *British Journal for the Philosophy of Science* 4, 147-157.
- Cherry, Colin: 1966, *On Human Communication*, Second edition, Cambridge, Massachusetts: MIT Press.
- Evans, R. D.: 1965, *The Atomic Nucleus*, New York: McGraw-Hill.
- Hempel, C. G.: 1965, 'Aspects of Scientific Explanation', in *Aspects of Scientific Explanation and Other Essays*, New York: Free Press.
- Hempel, C. G.: 1968, 'Maximal Specificity and Lawlikeness in Probabilistic Explanation', *Philosophy of Science* 35, 116-133.
- Railton, P.: 1978, 'A Deductive-Nomological Model of Probabilistic Explanation', *Philosophy of Science* 45, 206-226.