Probabilistic Explanation

Michael Strevens

In *Physical Theory: Method and Interpretation*,
Lawrence Sklar (ed.), Oxford University Press

1. Varieties of Probabilistic Explanation

Science turns to probabilistic, as opposed to deterministic, explanation for three reasons.

Most obviously, the process that produces the phenomenon to be explained may be irreducibly indeterministic, in which case no deterministic explanation of the phenomenon will be possible, even in principle. If, for example, the laws of quantum mechanics are both probabilistic and fundamental—as most scientists believe—then any explanation of, say, an episode of radioactive decay can at best cite a very high probability for the event (there being a minuscule probability that no atom will ever decay). The decay, then, must be explained probabilistically.

Because all the world's constituents conform to quantum dictates, it might seem that, for the very same reason, everything must be given a probabilistic explanation. For many phenomena involving large numbers of particles, however, the relevant probabilities tend to be so close to zero and one that the processes producing the phenomena take on a deterministic aspect. It is traditional in the philosophy of explanation to treat the corresponding explanations as deterministic.

Thus, you might think, there will be a simple division of labor between probabilistic and deterministic explanation: probabilistic explanation for phenomena involving or depending on the behavior of only a few fundamental-level particles, due to the indeterministic aspect of quantum mechanical laws, and deterministic explanation for higher-level phenomena where quantum probabilities effectively disappear. However, even high-level phenomena are routinely given probabilistic explanations, for reasons that have nothing to do with metaphysical fundamentals.

In some cases, the recourse to probability is for epistemic rather than metaphysical reasons. Although the phenomenon to be explained is produced in an effectively deterministic way, science's best model of the process may be missing some pieces, and so may not predict the phenomenon for sure. In such a case, the explanation is typically given a probabilistic form. Whatever the model says about the phenomenon is put in statistical terms—perhaps as a probability of the phenomenon's occurrence, or as a change in the probability of the phenomenon's occurrence brought about by certain factors—and these statistical facts are offered as a partial explanation of what has occurred.

There are many examples to be found in medicine. If a heavy smoker contracts emphysema, his or her smoking is typically cited as a part of the explanation of the disease. Smoking probabilifies emphysema, but we do not know enough about its etiology to see for sure whether any particular heavy smoker will become emphysemic. Thus our best explanation of a heavy smoker's emphysema must be probabilistic. Though we will perhaps one day be able to do better, we find the present-day probabilistic explanation enlightening: if it is not the best possible explanation of emphysema, it is certainly a fairly good explanation.

A third occasion for probabilistic explanation arises in certain cases where the process producing the phenomenon to be explained is rather complex and could have produced that very phenomenon in a number of different ways. In such cases, there appears to be explanatory value in a description of the process that abstracts away from the details that determine that the phenomenon occurred in the particular way that it did, and that presents only predisposing factors that make it highly likely that some such process would occur.

Perhaps the best examples are to be found in statistical physics. To explain why a gas rushes into a vacuum so as to equalize its density everywhere, you might recount the deterministic details in virtue of which each particle ends up in a state that, in the large, constitutes an equalization of density. You would have to cite the initial position and velocity of each particle, and derive its later position from these initial conditions using the appropriate laws of molecular dynamics.

Statistical physics tells the story in a different way. It can be shown (though the details are disputed) that almost every set of initial conditions leads to a gas's equalizing its density. The demonstration invokes only some quite general properties of molecular dynamics, so it is far simpler than the demonstration suggested in the previous paragraph, but because the *almost every* is not an *every*, it is not deterministic; statistical physics, as the name suggests, gives the explanation an explicitly probabilistic cast.

If the statistical explanation of density equalization were less satisfactory than the deterministic explanation, this case could be assimilated to the emphysema case, as a use of probability to fill in, for the time being, a gap left by scientific ignorance. But most writers (though certainly not all) would agree that the statistical explanation is superior to the deterministic explanation: you do not fully understand why gases fill vacuums until you see why almost any set of initial conditions leads to equalization, and once you appreciate this fact, seeing that the particular initial conditions of your particular gas led to equalization adds little or nothing of explanatory value. Here, then, probability is introduced into an explanation not because it is either metaphysically or epistemologically unavoidable, but because it enhances the explanation, providing more insight than the deterministic alternative. Even if we were able to construct a deterministic explanation of some gas's expansion, we would prefer the statistical explanation.

Another sort of explanation that perhaps belongs in the same class as the explanations of statistical physics is the probabilistic explanation of frequencies of outcomes obtained on simple gambling devices. To explain the fact that a large number of coin tosses, for example, turn up heads about one-half of the time, we typically cite the one-half probability of a tossed coin's landing heads, though we know that coin tossing is an effectively deterministic process. From the one-half probability we can derive a very high probability of a frequency of heads approximately equal to one-half, but without being able to make any predictions about the particular sequence of heads and tails that realizes the frequency. As in the case of the gas, we extract from our knowledge of the physics of coin tosses the fact that almost any set of initial conditions for a long series of tosses will produce a frequency of heads of about one-half (Strevens 1998), and we prefer the corresponding explanation to a deterministic explanation that begins with the exact initial conditions of the series, derives a particular sequence of heads and tails, and calculates the frequency of heads in the sequence.

A philosophical account of probabilistic explanation must do some kind of justice to the three varieties of probabilistic explanation I have described. It must show why probabilistic explanation is acceptable if the world is indeterministic; it must show why it is also acceptable in cases where our knowledge of the world precludes our constructing a deterministic explanatory model; and it must show why probabilistic explanation is preferable to deterministic explanation in statistical physics and certain other domains.

Then again, an account of probabilistic explanation might be in part a debunking of some or all of these claims, for example, an argument that scientists and philosophers are mistaken in thinking that probabilistic explanation is ever preferable to deterministic explanation—or even an argument that there can be no such thing as probabilistic explanation. Deflationary arguments of this sort can be found in the literature, but most theories of probabilistic explanation attempt to make sense of, rather than to deny the existence of, successful probabilistic explanation in quantum mechanics, medical science, and statistical physics.

Accounts of scientific explanation can be classified in two ways: according to their conception of the nature of explanation, or according to the formal criteria they impose on an explanation. These properties are of course linked, but not especially strongly. Two theorists might agree that an explanation should consist of a list of factors that are statistically relevant to the phenomenon to be explained (a formal criterion), yet disagree on the question why a statistically relevant factor casts light on the phenomenon. Or they might agree that explanation is ultimately concerned with giving the causes of a phenomenon, yet disagree on the question whether the causes should be presented in the form of a logical argument.

A treatment of approaches to explanation that follows the traditional divisions in the literature results, for this reason, in a certain amount of cross-classification: as you will see, some accounts of explanation might reasonably be placed under more than one of the organizational headings employed in this essay. When Hempel initiated the systematic study of probabilistic explanation in 1963, however, he clearly specified both a formal criterion for probabilistic explanation and a doctrine of the nature of the underlying explanatory relation. This provides a place to begin the survey of approaches to probabilistic explanation that is both especially firm and chronologically apt.

2. Nomic Expectability and the Inductive-Statistical Account

An explanation, Hempel writes,

> Shows that, given the particular circumstances and the laws in question, the occurrence of the phenomenon *was to be expected*; and it is in this sense that the explanation enables us to *understand why* the phenomenon occurred (Hempel 1965, 337).

Although this passage concerns deterministic explanation, Hempel's treatment of probabilistic explanation is also based on the posit that scientific understanding is a matter of nomically based expectation: a probabilistic explanation, Hempel proposes, uses a statistical rather than a deterministic law to show that the phenomenon to be explained was to be expected (Hempel 1965, §3).

From this conception of nomic expectability as the underlying explanatory relation, Hempel extracts a formal criterion for probabilistic explanation. Or rather, he extracts two criteria, one for the explanation of singular events and one for the explanation of laws. A probabilistic law can be explained by deriving the law deductively from other laws and (if necessary) initial conditions. This is precisely the same formal criterion that Hempel gives (tentatively; see Hempel and Oppenheim 1948, n33) for the explanation of a deterministic law.

The formal criterion for the probabilistic explanation of singular events constitutes, by contrast, a significant departure from the deterministic case. For this reason, and because the probabilistic explanation of singular events perhaps raises more questions than the probabilistic explanation of laws, philosophers writing on probabilistic explanation have focused on the explanation of events. This essay will, of necessity, reflect the bias towards event explanation.

An event explanation must show, according to Hempel's doctrine of nomic expectability, that the event to be explained was to have been expected. As a consequence, an explanation functions like, and for Hempel is identical to, an argument. If the event is produced deterministically, the argument can be deductive. If not, it can only be (if not trivial) inductive. Thus Hempel proposes that a probabilistic explanation of an event is a sound law-involving inductive argument to the effect that the event occurred. This he calls the inductive-statistical, or IS, account of probabilistic event explanation. The criteria for the soundness of an inductive argument differ in interesting ways from the criteria for the soundness of a deductive argument; many of the more controversial features of the IS account arise from the differences.

Suppose (Hempel's example) that you are assigned to explain why Jen Jones, who contracted a streptococcus infection, recovered within a week. It turns out that she was given penicillin, and that 90% of all strep patients recover quickly, that is, within a week, when given penicillin. You explain Jones' swift recovery probabilistically by citing its high probability, given the administration of penicillin.

According to Hempel, your explanation should be understood as an inductive argument with the event to be explained—that Jen Jones recovered within a week—as its conclusion, and the relevant probabilistic facts as its premises. Hempel formalizes the argument as follows:

> Jen Jones had a strep infection
> Jen Jones was given penicillin
> 90% of all strep patients recover within a week when treated with penicillin
>
> Jen Jones recovered within a week

Because the argument is not deductive, the premises do not entail the conclusion; rather, they bestow a 90% inductive probability on the conclusion. This is high enough, as Hempel writes, that you should have expected Jones to recover. Because your expectation is based in part on a statistical law, it confers scientific understanding.

Hempel's probabilistic event explanations differ from his deterministic event explanations in two ways, corresponding to two ways in which inductive logic differs from deductive logic. First is the matter of the probability itself. An inductive argument must confer a sufficiently high probability on its conclusion to justify an expectation that the conclusion holds. Typically, it is required that the probability be greater than 50%. It follows, on the IS account, that events with a probability of 50% or less cannot be explained. Since the probability of even a heavy smoker's contracting emphysema is less than 50%, for example, it seems that citing smoking alone cannot explain emphysema, though citing smoking together with other predisposing factors might be sufficiently probabilifying to be explanatory. That is a controversial claim.

Second, it is not enough, in order to provide a sound inductive argument that the recovery will occur, to find some collection of facts that together probabilify the recovery to a degree greater than 50%. The facts must in a certain sense be complete. To illustrate this point, Hempel asks you to imagine that Jones' infection is known to be resistant to penicillin. Suppose that the probability of swift recovery when infected with a penicillin-resistant strain is 10%, whether or not penicillin is administered. Then you ought not to have expected Jones to recover swiftly.

The fact of penicillin resistance in no way undermines the truth of the premises of the inductive argument spelled out above. It does undermine the argument itself; typically, it is said that an inductive argument is not sound unless it includes all relevant information in its premises. Because the argument for swift recovery omits the information concerning penicillin resistance, it does not provide inductive support for its conclusions. Thus, the power of an inductive argument depends not only on its internal structure and on the truth of its premises, but also on the stock of background knowledge.

For essentially the same reasons, Hempel claims that the argument does not, in the penicillin resistance case, explain its conclusion, a consequence reflected in actual explanatory practice. It is not possible for Hempel to require that an IS explanation cite all relevant background knowledge among its premises, since the fact that the explanandum occurred is normally known, and so would have to be included, making the argument deductive and, because not law-involving, unexplanatory. Other, related items of inductively relevant information must also be excluded as a matter of course from the explanatory argument, for example, evidence that the explanandum occurred. Hempel formulates what he calls the *requirement of maximal specificity* to find the right balance of inclusion and omission. The details will not be discussed here.

The requirement of maximal specificity demands that (almost) all relevant background knowledge be included in an inductive argument if it is to serve an explanatory function, but unknown facts that would be inductively relevant if known do not fall within the scope of the requirement. Thus if Jen Jones' infection is resistant to penicillin, but this fact is not known to the medical community, the argument above qualifies as a good explanation. Certainly it is reasonable to count the argument as a piece of sound inductive reasoning—inductive reasoning is always provisional—but it does not seem so reasonable to say that penicillin explains Jones' recovery provided that her doctors do not discover that hers is a resistant strain. The doctors may think that they are giving a good explanation, but they are wrong—so you would like to say. Hempel's account, however, has the consequence that they are genuinely giving a good explanation; only if they find out about the resistance does their explanation cease to confer understanding.

These, then, are the two principal objections to the IS account of event explanation (ignoring, for the sake of this essay, objections that apply equally to Hempel's deductive-nomological account of event explanation, such as the charge that an expectability account illegitimately allows effects to explain their causes): first, that the IS account allows only events with relatively high probabilities to be explained, whereas science considers itself able to explain low-probability events, and second, that the IS account relativizes probabilistic explanation to an epistemic background, where we find in scientific practice no such relativity.

The high probability objection has spawned a debate that has been conducted to some extent independently of the rest of the philosophy of probabilistic event explanation, a debate that concerns the extent to which the factors cited in an explanatory argument must probabilify the event to be explained, and whether factors that lower the probability ought to appear in an explanation (as on the IS account, you will note, they must). I defer substantive discussion of the debate to section 5.

The epistemic relativity objection has motivated a search for a conception of the explanatory relation, and a concomitant account of probabilistic explanation, according to which the probabilities involved in explanation are not inductive probabilities but something more objective, usually real physical probabilities.

3. Causal Approaches

Famous counterexamples to Hempel's deductive-nomological account of explanation, and more broadly to his conception of explanation as nomic expectability, such as the case of the length and period of a pendulum (see Woodward, *Explanation*, this volume), have by now convinced the great majority of philosophers that an account of explanation must provide a starring, if not exclusive, role for causation. The simplest way to do so is to hold that what explains an event are its causes and the background conditions and laws (or causal generalizations) in virtue of which they are causes.

The causal conception has naturally been transplanted to the territory of probabilistic event explanation, but it is not clear exactly where it should take root. Most commonly, it is suggested that physical probabilities (sometimes called *chances*) can be understood as a kind of causal disposition. A probability explains the outcome that it is the probability of, on this view, the same way that any disposition explains the event that it is a disposition to produce—in the same way, for example, that a magnet's disposition to attract any nearby iron filings explains a clump of filings' jumping and clinging to the magnet. Among the dispositionalists may be counted Coffa (1974), Fetzer (1974), and as godfather if not participant, Jeffrey (1969).

Whereas in the IS model, it is inductive probabilities that explain, in a dispositional theory, it is physical probabilities, interpreted as a kind of cause. This gives the dispositionalist a more satisfying account of the case of penicillin resistance, discussed in the previous section, than Hempel is able to offer. If

Jen Jones' infection is not known to be penicillin-resistant, then her treatment with penicillin confers, uncontroversially, a high inductive probability on her swift recovery. The IS account is committed, controversially, to counting this probability as explanatory. The dispositionalist is not: they require that the penicillin, if it is to contribute to the explanation, should contribute to a physical disposition to recovery; in the case where the infection is resistant, it clearly does not do so.

A dispositional view of probability and probabilistic explanation can also avoid the worst consequences of the IS account's requirement that an event have a high probability if it is to be probabilistically explained, in either of two ways. First, an explanation may be allowed to cite less than the full complement of causes (though presumably, more causes are on the whole better). If you explain a case of emphysema by pointing to heavy smoking, you indicate one cause of the emphysema—one part of the basis for the probability of the smoker's contracting, thus by hypothesis the smoker's disposition to contract, emphysema—and so give a genuine, if suboptimal, explanation of the emphysema. On the IS account, by contrast, your attempt fails to count as an explanation at all.

Second, in the case where an event's only cause fails to probabilify it highly, on a causal account the citation of that cause, and thus of the low probability, nevertheless constitutes a causally exhaustive hence an optimal explanation of the event. Suppose, for example, that the probability of a certain atom's undergoing radioactive decay in a given time period is low. On the IS account, the event of the decay's occurring within the time period, if it happens, cannot be explained. On a causal approach it can be given a complete explanation by delineating the relevant causal information, that is, the aspects of the structure of the atom and the laws of quantum mechanics that fix the probability of decay.

Many commentators have found the idea of a probabilistic disposition obscure. Whereas the familiar, deterministic dispositions such as fragility and paramagnetism can be given a counterfactual analysis, any parallel analysis of probabilistic dispositions seems itself to contain an ineliminable reference to probability.

It is not necessary to understand probabilities as dispositions in order to take a causal approach to probabilistic explanation. One alternative, suggested by Paul Humphreys, takes the factors that affect the value of an event's probability, but not the probability itself, as the causes of the event. Heavy smoking is a cause of emphysema, then, because it increases the probability of emphysema, but the probability itself is not a causal disposition. Indeed—taking an extremely deflationary attitude to probabilities—Humphreys declares them to be "literally nothing" (Humphreys 1989).

A different approach, which is compatible with either Humphrey's deflationism or with dispositionalism, is to hold that when probabilities enter into causal explanation, they do so in the guise of probabilistic causal laws. Heavy smoking is a cause of emphysema, on this view, because there is a causal law about emphysema that connects it probabilistically to smoking. Thus probabilities are a part of causal explanation because the causes do their causing probabilistically; it is left open whether probabilities are themselves causal. To the philosopher of explanation, this view has the advantage of leaving some of the more contentious issues to the metaphysicians. In practice, however, it has not been so easy for explanatory causalists to avoid the metaphysics of probability (this writer included).

4. Probabilistic Relevance Accounts

What probabilistic relevance accounts have in common is not a conception of the explanatory relation—as expectability, causality, or something else—but an aspect of the formal criterion they offer for a good explanation. They are agreed on two things. First, they hold that a probabilistic explanation is either a list of factors that are probabilistically relevant to the event to be explained, or some other structure designed to exhibit such factors and their relevance, such as a deductive argument. Second, they take this fact about the formal criterion to be a more fundamental and reliable datum about explanation than any intuitions you might have about the underlying nature of scientific understanding. They differ, then, from an account such as Humphreys', which also requires that an explanation cite probabilistically relevant factors, by leaving open the question of the nature of the explanatory relation.

Hempel's formal criteria for both deterministic and probabilistic explanation contain no effective safeguard against the inclusion, in an explanatory argument, of irrelevant information, that is, in the case of event explanation, information that plays no role in entailing the occurrence of the event to be explained (in the deterministic case) or probabilifying its occurrence (in the probabilistic case). Its appearance in an explanatory argument implies that such information is explanatorily relevant, but intuition suggests otherwise. The expectability approach could be insulated against this objection in the following way: an explanation is a list of factors that, first, provide good reason to believe that the phenomenon to be explained occurs, and second, are statistically relevant to its occurrence.

Wesley Salmon suggests instead a radical departure from Hempel's argument-centered formal criterion for explanation: all that matters for an explanation is the satisfaction of the second requirement, relevance to the event's occurrence. Further, Salmon suggests, the relevance relation in question is not epistemic but physical: a factor is relevant to the occurrence of an event if it affects the physical probability of the event (Salmon 1970).

For explanatory purposes, a factor d is not considered to affect the probability of an event e if d is screened off from e by a further event c but not vice versa, meaning that d does not affect the probability of e in the presence of c, but c does affect the probability of e in the presence of d. Thus, for example, d is not said to affect the probability of e in the case where c is a common cause of d and e, despite the fact that d and e will typically be correlated.

According to what Salmon calls the statistical relevance account, an event explanation is a list or table of all factors that make a difference to the physical probability of an event's occurring. Salmon takes an expansive view of this table of relevance: he suggests that it should include information about factors that were relevant to unrealized alternatives to the event to be explained, and factors that were not present but that would have been relevant if present. Factors that are negatively relevant are also considered explanatory; that is, factors that lower the probability of an event's occurrence should be included in the event's explanation. Salmon defends these requirements only in passing; they may be regarded as optional elements of the statistical relevance account.

Peter Railton has also offered what may be considered a statistical relevance account of probabilistic event explanation, though he calls his theory the deductive-nomological-probabilistic account (Railton 1978). Railton proposes than an event is explained by deducing its physical probability from the relevant laws and initial conditions. In contrast to the IS account, the explanatory argument is deductive, and has the physical probability of the event as its conclusion, rather than the event itself. Further, the argument is explicitly required to contain only facts essential to the deduction. A Railtonian explanation, then, will contain only facts found in a Salmon explanation, arranged in the form of a deduction. It will not contain all the facts that Salmon requires, however: it will mention only factors that were present and that contributed to the probability of the event to be explained (as opposed to its unrealized alternatives).

Salmon and Railton do not link their accounts to a particular conception of the explanatory relation. Unofficially, Salmon perhaps takes statistical relevance itself as the relation, at least in his early work, and Railton occasionally talks in a causal idiom. But officially their accounts are open to many interpretations, for example, to either dispositionalism or its denial.

Neither account, however, is compatible with Hempel's expectability conception of the explanatory relation, since both allow the probability of the explained event to be as low as you like. Provided that you can cite all the factors that play a part in determining the probability of an event, you can explain the event, even if it is very unlikely.

Achinstein (1983) presents the following counterexample to any probabilistic relevance account of event explanation (including theories such as Humphreys'). Petra takes poison. This particular poison has a 90% chance of killing anyone who takes it within 24 hours. As it happens, Petra survives the poisoning, but is run over by a bus exactly 24 hours after taking the poison. The poison probabilifies her death but does not explain it. (Gluck and Gimbel (1997) offer a more sophisticated version of this argument.) Thus there is more to explanation than probabilistic relevance.

In reply, a proponent of probabilistic relevance might argue that there must have been something about Petra—her high metabolism or her having breakfasted on the antidote—that prevented her dying. In the context of this intervening factor, the poison did not raise the probability of death after all, thus the probabilistic relevance approach does not count it as explanatorily relevant. Surely, however, the action of the poison might be genuinely indeterministic. Or it might be deterministic, but of a piece with the processes that are the subject matter of statistical physics. Either way, you must conclude that the probability of death was raised, but nothing came of it, and so that the probability-raiser is as a consequence no part of the explanation of death. The same problem arises for accounts of probabilistic causation (Menzies 1996); for a solution, see Strevens (2008), §11.3.

5. Elitism and Egalitarianism

Elitism is the blanket term I give to a preference for high probability in explanation; *egalitarianism* is a contrary indifference to probabilistic magnitude, and in the case of probability changes, perhaps even to sign. There are two questions in particular towards which elitist and egalitarian attitudes can be distinguished in event explanation: the question of the size of the probability attached to the event to be explained, and the question of the change in the probability of the event brought about by a statistically relevant factor. I call these respectively the *size* debate and the *change* debate.

Three main positions have been taken in the size debate. The first is Hempel's: an explanation must show that the event to be explained has a high probability (at least greater than one-half); thus, only high-probability events can be explained. I call this view *extreme elitism*. It has the consequence that heavy smoking alone cannot explain emphysema. Hempel would stand by this conclusion even if smoking were the only way to contract emphysema (compare Scriven's (1959) famous example of paresis), claiming either that some other predisposing cause of emphysema must be cited, or if there is none, that we simply do not understand why the patient contracted emphysema.

The contrary view, that events with low probabilities can be explained, has two variants, one elitist and one egalitarian. According to the *moderate elitist*, the higher the probability of an event, the better it is explained by citing that probability. Low probability events can be explained, then, but not as well as high-probability events. According to the *egalitarian*, events are equally well explained regardless of their probability.

Most philosophers are agreed that there are explanations of low probability events to be found in science, and so that extreme elitism does not capture actual explanatory practice (though it might, of course, be regarded as a reformist proposal). It is more difficult to use scientific practice to adjudicate between moderate elitism and egalitarianism, since these views concur as to which events can be explained, disagreeing only as to how well they are explained. Strevens (2000) argues that the history of the development of statistical physics favors the moderate elitist position.

There are also more philosophical, which is to say a priori, considerations that can be brought to bear on the debate. Some writers hold that to explain an event is to show why it occurred rather than not occurring; they see a tension between this doctrine and the egalitarian view that both the occurrence and the non-occurrence of an indeterministically produced event can be explained equally well. (Salmon (1990), pp. 178–179 formulates, without endorsing, the argument.)

Other writers hold that an explanation that enumerates accurately all the causes of the event it explains is perfect. A low-probability event is perfectly well explained, then, if the factors that determine its low probability are its only causes, which is presumably the case if (and only if) the low probability is of the irreducible variety, in particular, if it is a quantum mechanical probability.

The change debate concerns the relationship between a factor's probabilistic impact on the event to be explained and its explanatory importance. Are factors that make a larger probabilistic impact explanatorily more important? Are factors that make a negative impact explanatorily important? As in the size debate, three positions can be distinguished. (Though I give the positions in the size and change debates similar names, you will see that they are distinct.)

On egalitarian views, the size of the impact does not affect a factor's explanatory value. Citation of a factor that makes a large difference to the probability of the event to be explained adds no more to an explanation than citation of a factor that makes only a small difference. To put it another way, appreciating the probabilistic role of a factor that makes a large difference no more illuminates the occurrence of the event to be explained than does appreciating the probabilistic role of a factor that makes virtually no difference. If heavy smoking quintuples your probability of contracting emphysema, while having a wood-burning fire increases the probability by 1%, then an explanation of a person's emphysema that cites only their wood-burning fire is just as illuminating as an explanation that cites only their smoking. Despite the awkward sound of these consequences, many writers on probabilistic explanation, including the founder of the statistical relevance approach, Wesley Salmon, tend to egalitarianism.

A question that divides egalitarians concerns factors that lower the probability of the event to be explained. *Moderate egalitarians* hold that such factors are not explanatory; *extreme egalitarians* that they are. Salmon has argued most explicitly for extreme egalitarianism, though even he forbids explanations that cite only probability-lowering factors (Salmon 1984, 46). Humphreys is another staunch extreme egalitarian. Observe that any formal criterion for probabilistic explanation that requires the presentation of all factors that play a role in fixing the probability of the explanandum will render probability-lowerers explanatorily relevant. Railton's and even Hempel's accounts of event explanation fit this description.

On the other side of the divide lies *moderate elitism*, the view that the more a factor increases the probability of the event to be explained, the more it contributes to the event's explanation.

Although the resolution of the change debate would cast much light on the nature of probabilistic explanation—not least the problem of explanatorily irrelevant probability-raisers, such as Petra's poisoning in the previous section—philosophers have not gone much further than laying out the issues. Such arguments as there are tend to reflect the arguments in the size debate; for example, on the causal approach to explanation you might think (as Humphreys argues) that a factor that decreases the probability of an event is a part of the causal history of that event, and so qualifies alongside the event's probability-raisers for a place in the explanation of the event. You will then tend to extreme egalitarianism in both the size and change debates.

Or if you are a dispositionalist, you might think that the probability of an event quantifies the "force" bringing the event about. The more force, the easier it is to understand the fact that the event occurred. High probability events are therefore better explained, and greater contributions to the probability of an event are greater contributions to its explanation. Result: moderate elitism in the size and change debates. Although taking a position in the size debate does not force you to take any particular position in the change debate, then, there is an affinity to be found between same-named positions.

6. Probability and Determinism

Of the three varieties of probabilistic explanation described in section 1, only one paradigmatically involves the citation of an irreducible probability, such as a probability stipulated by the laws of quantum mechanics. The other two—explanation in which probability fills an epistemic gap, as in the case of emphysema, and explanation in complex systems where a probabilistic theory captures predisposing causes while abstracting from the details of particular initial conditions, as in statistical physics—work perfectly well even in deterministic systems.

A number of writers have cast doubt on the validity of probabilistic explanation in deterministic systems by way of the following two premises:

1. There are no physical probabilities in deterministic systems, and
2. Where there are no physical probabilities, there can be no probabilistic explanation.

This is an incredible conclusion; after all, explanation in statistical physics and other high-level sciences that employ probability in a similar way is a well-established fact of scientific life, hardly ripe for overthrow by philosophers. The premises, then, merit a closer look.

The view that there can be no physical probabilities in a deterministic system has a certain intuitive appeal (Schaffer 2007). However, of the various philosophical accounts of the nature of physical probability, several allow the existence of such probabilities: the various versions of the frequency account, Popper's propensity theory, and proposals to find the basis of certain systems' physical probabilities in the mathematical properties of the systems' dynamics (Hopf 1934; Strevens 2003, 2011; Abrams 2012).

Salmon, for example, adopts a version of the frequency theory to provide the metaphysical foundations for his statistical relevance account. But frequentism has fallen out of favor, and the view that probabilities can only be some kind of irreducible propensity has taken its place in the literature as the default view. This is the one interpretation of probability that disallows physical probability in deterministic systems. An apparent impasse, then: either the dominant metaphysics of probability must be discarded or perhaps augmented, or the practice of probabilistic explanation eviscerated.

There is a third way, however: premise (2) above, that there can be no probabilistic explanation without physical probability, is not as obviously true as it might seem.

One way around the premise is via accounts of probabilistic explanation that do not call directly on physical probability. Hempel's IS account is the preeminent example: what Hempel requires for explanation is in the first instance inductive, not physical, probability. (Hempel himself writes that the probabilistic laws in IS explanation must be based in real physical probabilities; however, he clearly has in mind a metaphysics of physical probability—presumably empiricist—on which there are probabilities wherever there are robust statistics.)

Not many contemporary philosophers would be willing to turn to the IS account to make room for probabilistic explanation in deterministic systems, however. There is a second option. It is generally accepted that, although the fundamental laws of nature are quite possibly indeterministic, deterministic explanation is legitimate when the fundamental probabilities are all close enough to zero and one. That is, deterministic explanation is possible, in special circumstances, in indeterministic systems. Perhaps, then, probabilistic explanation is possible, in special circumstances, in deterministic (or near deterministic) systems. More exactly, perhaps the apparatus of probabilistic explanation is especially well suited to capturing the explanatorily relevant aspects of certain systems that are at root deterministic, and so that, by some philosophers' lights, contain no physical probabilities.

Railton (1981) suggests that the probabilistic element of statistical physics functions in explanation to capture the robustness of the underlying physical processes. Since I framed the presentation, in section 1, of probabilistic explanation in statistical physics in just these terms, Railton's view will sound familiar. The following discussion will flesh out the view, taking as a framework something somewhat stronger than Railton's suggestion. The framework is comprised of the following four posits (of which the first can be regarded, for the sake of the philosophical dispute, as uncontroversial):

1. In the domain of statistical physics (and, it might be added, many other areas), the kind of event you want to explain, such as the expansion of a gas to fill a vacuum, occurs given almost any set of initial conditions. Call this property *robustness*.

2. To understand a robustly produced event, you must grasp the reasons for the robustness of the underlying process.

3. To understand a robustly produced event, you need not appreciate the exact initial conditions that led to the event.

4. A probabilistic representation is especially good at capturing the explanatorily essential facts about robustness that are the subject of premise (2) while omitting the explanatorily uninteresting facts about the details of initial conditions that are the subject of premise (3).

Railton himself asserts (2), though without giving an account of the explanatory value of robustness; he would probably not endorse (3), and might or might not agree with (4). The following discussion explores more recent literature on robustness that better fits the framework.

The framework takes no position on the question of the existence of physical probabilities in systems to which its posits apply. You might believe that the robust elements of the systems' dynamics picked out by the probabilistic representation in some sense constitute genuine physical probabilities, you might deny it, or you might withhold judgment. The framework allows all three responses. A division of labor is therefore proposed: let philosophers of explanation articulate the explanatory value of probabilistic theories in systems with robust dynamics; let metaphysicians of probability decide whether the basis of such explanations includes genuine probabilities.

Event *e* is produced robustly. Why is understanding the basis of robustness explanatorily more important than understanding the exact initial conditions that led to *e*?

Jackson and Pettit (1992) argue for two modes of explanation, one in which the initial conditions are paramount and one in which the fact of robustness is what matters. In the latter mode, the aim of explanation is to discover similarities between the causal process that actually produced *e* and the causal processes that produce *e* in nearby possible worlds, that is, roughly, the causal processes that might have produced *e* if the actual process had not. In a robust system, these would-be producers of *e* are processes in the same system with different initial conditions. A good explanation, in Jackson and Pettit's second mode, points to what the processes have in common and overlooks their differences; it therefore cites the properties responsible for the system's robustness and ignores the particularity of initial conditions.

Woodward (2003) sees an explanation of *e* as a body of information, in the form of a causal model, as to how to causally manipulate the occurrence of events like *e*. Information about robustness is useful for manipulation, but information as to exact initial conditions is not, since changing the initial conditions of a robust process is unlikely to change the outcome of the process.

Strevens (2004) proposes that an explanation of *e* specifies (some of) the factors that made a difference to the causal production of *e*. A causal detail fails to make a difference to *e* if a causal model for *e* that abstracts away from its presence—that does not specify how things are with respect to that detail—is sufficient to entail *e*, at least with a high probability. (Note that the entailment in question must mirror a real causal process.) When *e* is produced robustly, then, the details of the initial conditions that produced *e* are not difference-makers, but the properties responsible for the robustness are. Unlike Jackson and Pettit and Woodward, then, Strevens holds that the explanatory relevance of robustness and the irrelevance of initial conditions are due to the role that these factors play in the actual causation of *e*.

There are a number of ways, as you can see, to account for the explanatory superiority of a probabilistic model that cites robustness and ignores initial conditions over a strictly deterministic model that cites the exact initial conditions. But what role does probability itself play in a robustness-based model?

One answer is that, in the case where only *almost all* initial conditions lead to the event *e* to be explained, and where the conditions that do not lead to *e* are thoroughly mixed in with those that do, a deterministic model is for all practical purposes impossible. To specify the initial conditions sufficiently finely to rule out any possibility of conditions not producing *e*, you would have to disgorge such torrents of detail that you might as well simply state the actual initial conditions, defeating the purpose of a robustness-based model.

Another, more interesting answer is that to understand the basis of robustness itself, you must make use of the conceptual inventory of probability theory. This view is suggested by Strevens (2003, 2008), in which it is argued that the robustness-like properties of the processes described by statistical physics, evolutionary theory, and perhaps some of the social sciences are best understood by invoking all the apparatus of probability theory—probability densities, stochastic independence, the law of large numbers, the central limit theorem—regardless of whether there are "metaphysically real" physical probabilities at work in the system. On Strevens' proposal, an explanatory model for *e* that cites only the robustness of the process leading to *e* might be fully deterministic, yet it might also be probabilistic in the sense that it uses probabilistic concepts to derive and so to explain the robustness. The value of probabilistic explanation in deterministic systems is, then, not merely that probabilistic thinking is unavoidable, but that it is invaluable, in some cases even where determinism reigns.

References

Abrams, M. (2012). Mechanistic probability. *Synthese* 187:343–375.

Achinstein, P. (1983). *The Nature of Explanation*. Oxford University Press, Oxford.

Coffa, A. (1974). Hempel's ambiguity. *Synthese* 28:141–163.

Fetzer, J. (1974). A single case propensity theory of explanation. *Synthese* 28:171–198.

Gluck, S. and S. Gimbel. (1997). An intervening cause counterexample to Railton's DNP model of explanation. *Philosophy of Science* 64:692–697.

Hempel, C. G. (1965). Aspects of scientific explanation. In *Aspects of Scientific Explanation*, chap. 12, pp. 331–496. Free Press, New York.

Hempel, C. G. and P. Oppenheim. (1948). Studies in the logic of explanation. *Philosophy of Science* 15:135–175.

Hopf, E. (1934). On causality, statistics and probability. *Journal of Mathematics and Physics* 13:51–102.

Humphreys, P. (1989). *The Chances of Explanation*. Princeton University Press, Princeton, NJ.

Jackson, F. and P. Pettit. (1992). In defense of explanatory ecumenism. *Economics and Philosophy* 8:1–21.

Jeffrey, R. C. (1969). Statistical explanation vs. statistical inference. In N. Rescher (ed.), *Essays in Honor of Carl G. Hempel*, pp. 104–113. D. Reidel, Dordrecht. Reprinted in Salmon et al. (1971).

Menzies, P. (1996). Probabilistic causation and the preemption problem. *Mind* 105:85–117.

Railton, P. (1978). A deductive-nomological model of probabilistic explanation. *Philosophy of Science* 45:206–226.

———. (1981). Probability, explanation, and information. *Synthese* 48:233–256.

Salmon, W. C. (1970). Statistical explanation. In R. G. Colodny (ed.), *The Nature and Function of Scientific Theories*, pp. 173–231. University of Pittsburgh Press, Pittsburgh. Reprinted in Salmon et al. (1971).

———. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton University Press, Princeton, NJ.

———. (1990). *Four Decades of Scientific Explanation*. University of Minnesota Press, Minneapolis.

Salmon, W. C., R. Jeffrey, and J. Greeno. (1971). *Statistical Explanation and Statistical Relevance*. University of Pittsburgh Press, Pittsburgh.

Schaffer, J. (2007). Deterministic chance? *British Journal for the Philosophy of Science* 58:113–140.

Scriven, M. (1959). Explanation and prediction in evolutionary theory. *Science* 30:477–482.

Strevens, M. (1998). Inferring probabilities from symmetries. *Noûs* 32:231–246.

———. (2000). Do large probabilities explain better? *Philosophy of Science* 67:366–390.

———. (2003). *Bigger than Chaos: Understanding Complexity through Probability*. Harvard University Press, Cambridge, MA.

———. (2004). The causal and unification approaches to explanation unified—causally. *Noûs* 38:154–176.

———. (2008). *Depth: An Account of Scientific Explanation.* Harvard University Press, Cambridge, MA.

———. (2011). Probability out of determinism. In C. Beisbart and S. Hartmann (eds.), *Probabilities In Physics.* Oxford University Press, Oxford.

Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, Oxford.