



Universiteit
Leiden
The Netherlands

Explanation and determination

Gijsbers, V.A.

Citation

Gijsbers, V. A. (2011, August 28). *Explanation and determination*. Retrieved from <https://hdl.handle.net/1887/17879>

Version: Corrected Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/17879>

Note: To cite this publication please use the final published version (if applicable).

Chapter 7

Indeterministic Explanation I

In this chapter and the next we look at apparently non-deterministic explanations of events that are produced by deterministic processes. In the current chapter, I show how explanations that would traditionally be interpreted as examples of the inductive-statistical (IS) model or of the statistical relevance model can be interpreted as examples of the determination theory. For this purpose I construct a model, the non-specific deductive or ND model, that describes the form seemingly inductive-statistical explanation take within the determination theory. I argue that my theory and the ND model cannot only handle such explanations, but can solve important philosophical problems concerning the role of probability in explanation – this is the burden of sections 7.2 and 7.3. The ND model naturally leads to the question of degrees of explanatory strength, which I discuss as an aside in section 7.4.

In sections 7.5 and 7.6, I turn to the explanation of truly indeterministic events. On my theory, such explanations do not exist; but I hope to show that this does not violate our intuitions (although some people have thought it does), and that it is not an affront to modern physics.

In chapter 8, we will go over some of the same ground. There I tackle the recent argument of Michael Strevens that determined events are often best explained in a non-deterministic way. Because of the length and technicality of this discussion, it seemed better to split it off into its own chapter. Neither chapter presupposes the other.

7.1 The non-specific deductive model

The determination theory claims that all explanation is a matter of determination, and this immediately raises the problem of how to deal with putative examples of non-deterministic explanation. These examples can be split into

two classes: explanations of events that were produced by truly indeterministic processes, and non-deterministic explanations of deterministically produced events. I will discuss the first class in sections 7.5 and 7.6, and focus on the second class now (and again in chapter 8).

Throughout this discussion, some of the things that the determination theory considers a necessary part of all explanations will remain implicit. In particular, we will often not write down the contrast classes if they are obvious, we will silently assume that the interventionist condition 4 has been met, and we will never spell out that the proposition in the explanandum is true. Not doing so would prove very tedious.

Let us look at the following famous example of an explanation (Scriven 1959 [121]):

- (1) Jones got paresis because he had untreated syphilis.

This explanation has been much discussed because, although untreated syphilis *can* lead to paresis, it does so only in a minority of cases. There are presumably causal factors that determine whether or not any given person who has untreated syphilis will also get paresis; but these factors are unknown to current medical science. Thus, explanation (1) is the best explanation we can give of Jones's paresis, even though the explanans does not imply the explanandum with certainty (and perhaps not even with a probability larger than 0.5). According to Salmon 1998 [113] (pp. 39-40), explanation (1) gives "some understanding of what happened and why"; it is an example of "inductive explanation" using a "statistical law", rather than of "deductive explanation" using a "universal law".

Salmon (1998 [113], p. 40) suggests that explanation (1) is to be understood as an instance of what is essentially Hempel's **inductive-statistical model** (Hempel 1965 [41]; Salmon 1989 [112], pp. 53-58):

$$\begin{array}{l}
 F(a). \\
 F(a) \implies P(G(a)) = p. \\
 \hline
 G(a).
 \end{array}
 \quad [p]$$

If this is the correct reconstruction of (1), and if explanation (1) gives us some understanding of the explanandum, then it cannot be the case that determination is a necessary component of all explanation. The IS model is not compatible with my theory.

I would now like to suggest that it is *not* the correct reconstruction of (1). In this section, I will give an alternative reconstruction that is compatible with the determination theory; in the two sections that follow, I will argue that my alternative allows us to answer some of the pressing questions concerning explanation and probability, and is therefore preferable to the IS model.

Remember that we assume that Jones's paresis has been deterministically produced. This assumption can be formalised as follows, where a is Jones, F is the property of having untreated syphilis and G is the property of having paresis:

$$\exists H : (F(a) \wedge H(a) \implies G(a)) \wedge H(a).^1$$

That is, there is some property H such that H and the untreated syphilis together always lead to paresis, and a has this property.

Now, in the case of paresis the exact nature of H is unknown; but insofar as we are justified to believe that the paresis is the result of a deterministic process, we are justified to believe that there is such an H and that it obtained in the case of Jones. This means we can see (1) as an instantiation of what I will call the **non-specific deductive model**, or the ND model, which gives the logical form of a certain class of explanations allowed by the determination theory:²

$$\begin{array}{l} F(a). \\ \exists H : (F(a) \wedge H(a) \implies G(a)) \wedge H(a). \\ \hline G(a). \end{array}$$

The logical form defined here is deductive: the fact that the unfortunate Jones has both untreated syphilis and property H makes it certain that he develops paresis. And again, although we may not know what H is, we can have good evidence that such a property H exists (namely, the evidence we have that the process in question is deterministic) and that Jones has this property (because he develops paresis). Thus, we may be justified in believing both premises of the argument – and indeed, anyone who believes that paresis is deterministically produced, and that syphilis is a necessary

¹We need something like second-order logic here, in order for the quantifier to range over properties; but going into more logical detail is not worth the trouble.

²As indicated above, we leave unsaid everything about contrast classes, interventions, and so on.

cause of paresis, *must* believe them. There is no epistemic argument against using the ND model.

Thus, we can accommodate the paresis example within the determination theory. This counters the counterexample, but it is not yet a positive argument in favour of the theory. For that, I will now spell out an intriguing (and tempting) result of my proposal: it dispenses with the discussion between what Strevens (2008 [132], p. 347) calls elitism and egalitarianism. Let us review this debate and see how the determination theory coupled with the ND model manages not just to avoid the problems, but even reconciles the conflicting intuitions in this debate. This is certainly a positive argument in its favour.

7.2 Probability: the debate

Strevens (2008 [132], p. 347) distinguishes two related debates about explanation and probability. The first debate is the *size debate*, and it is about the relation between explanatory power and the value of p in the IS model. Must p be bigger than some threshold, perhaps 0.5, for an instance of the inductive-statistical model to be a real explanation? If there is no such threshold, is it the case that all probabilities explain equally well, or do high probabilities explain better than low probabilities?

The second debate is the *change debate*, and it is about the explanatory power of probability-increasing and (especially) probability-decreasing causal factors. Is the explanatory power of a causal factor a function of the change in probability of the outcome due to this factor? Can causal factors that decrease the probability of the outcome be explanatory, or are only factors that increase the probability explanatory?

The distinction between the two debates has not always been made in the literature, but it will nevertheless be useful to discuss them separately.

We will review the **size debate** first. Hempel 1962 [40] and 1965 [41], where the IS model was first worked out, requires that p is high, so that the explanandum is “almost certain”. However, Hempel 1977 [43] relinquished this requirement in order to cope with counterexamples.³ One of these counterexamples (originally from Jeffrey 1969 [53], p. 110) is described in Salmon 1984 [111]:

If two heterozygous brown-eyed parents produce a brown-eyed child, that fact can presumably be explained statistically on the

³Hempel 1977 [43] is out of print and hard to track down; I base myself on remarks in Salmon (1984 [111], p. 140; 1998 [113], p. 294) and Smith (1990 [126]).

basis of the 0.75 probability of such an occurrence. If these same parents produce a blue-eyed child, that fact seems unexplainable because of its low probability. Nevertheless, as Richard Jeffrey (1969) and others have argued persuasively, we understand each of these occurrences equally well. To say that we can explain the one, but not the other, is strangely asymmetrical. (p. 294.)

According to Salmon's, Jeffrey's and Hempel's egalitarianism, then, the value of p is irrelevant to how good a statistical-inductive explanation is. This view has been dominant ever since, but there have been a few detractors. Mellor 1976 [80] urges that p must be high, although he rejects Hempel's original argument, which was based on the idea that explanation is a kind of inference. Instead, Mellor argues that explanations are meant to close or at least narrow the gap between what is and what must be: we ask for an explanation of E only when we know that E is the case, but do not see why it had to be the case. Obviously, an inductive-statistical explanation with a low p fails to show why E had to be the case. Mellor goes on to say that Jeffrey and Salmon have:

... a rather startling view of explanation. Smoking hardly seems to explain not getting cancer as well as getting it; still less does it seem to explain not getting it just *because* it would explain getting it. (Mellor 1976 [80], p. 238.)

Strevens 2000 [128] notes that egalitarianism implies that statistical mechanics gives equally good explanations of those cases where gas released from a small container expands to fill the entire room and of those cases where the gas remains in a tiny portion of the space accessible to it. He argues that we cannot accept this, since statistical mechanics was deemed a successful scientific theory *because* it ascribed high probability to those processes that we actually observe (the entropy-increasing ones). Egalitarianism thus "amounts to the claim that Maxwell and Boltzmann enhanced SM's explanatory power not one degree" (p. 374). Or suppose that heat always flows from cold to hot objects; then the egalitarianist has to claim that statistical mechanics furnishes a good explanation of this fact, even though it assigns a ridiculously low probability to it. Strevens concludes by arguing that higher probabilities explain better:

The paretic case shows that high probabilities are not necessary for explanation. But the nature of probabilistic explanation in statistical mechanics shows that high probabilities explain better than low probabilities. (p. 389.)

Thus, we have two conflicting intuitions about whether or not (and how well) we can explain events that have only a low probability given the causal factors we put into our explanation. In the next subsection, I will show how my theory deals with these examples and manages to reconcile the intuitions.

The **change debate** was started by Salmon 1971 [108], which claimed that even probability-decreasing factors could be used to explain events. Cartwright 1979 [15] disagreed:

It is true that spraying with defoliant causes death in plants, but it is not true that spraying also causes survival. Holding fixed other causes of death, spraying with my defoliant will increase the probability of the plant's dying; but holding fixed other causes of survival, spraying with that defoliant will decrease, not increase the chances of the plant's surviving. (Cartwright 1979 [15], pp. 425-426.)

Cartwright concludes that spraying cannot explain survival. If the plant survives it is despite the spraying, not because of it.

Rogers 1981 [105] attempts to defend Salmon by making the surprising claim that although the spraying is not a cause of the survival, it is nevertheless part of the causal explanation:

The defoliant may not cause an individual poison ivy plant to survive, but the causal explanation of its survival, given treatment with defoliant, is that it was sprayed and its probability of survival was .1. (Rogers 1981 [105], p. 216.)

Humphreys 1981 [50], Salmon 1984 ([110], p. 46) and Humphreys 1989 ([51], p. 100-101) clarify Rogers's idea. They claim that there are two kinds of causes – contributing causes, which increase the probability of the explanandum, and counteracting causes, which decrease the probability – and that both can feature in a good explanation, as long as at least one contributing cause gets mentioned.

According to Strevens 2008 [132] p. 352, “Humphrey's bold thesis” is that counteracting causes have explanatory relevance “due to their decreasing, rather than increasing, the probability of the explanandum.”. Strevens disagrees with this bold thesis. But it is perhaps an overstatement of Humphrey's ideas, since he writes that (p. 101): “[The list of counteracting causes] is not a part of the explanation of *Y* proper. The role it plays is to give us a clearer notion of how the members of [the list of contributing causes] brought about *Y* – whether they did so unopposed, or whether they had to overcome causal opposition in doing so.”

It is not entirely clear, then, whether Humphreys and Strevens disagree; and we do not need to decide the issue here. The moral *I* wish to take from the change debate is that, on the one hand, counteracting causes can hardly be said to explain the event in question (Cartwright, Strevens); but on the other hand, we often *do* mention them in explanations, and we may even believe that we understand what happened better when we do (Salmon, Humphreys). It is desirable that a theory of explanation makes sense of both these intuitions.

7.3 Probability: answers

What can we say about the size debate and the change debate from the perspective of the determination theory? Let us first notice that the questions as asked do not arise in my theory: unlike the inductive-statistical model, the non-specific deductive model does not feature a variable probability p . If my ND reconstruction of the syphilis explanation is correct, it does not matter whether untreated syphilis causes paresis in 99 percent or in 1 percent of the cases: what matters is whether it caused paresis in the case of Jones.

This result is, I think, intuitively right. Suppose that syphilis causes paresis if and only if the patient smokes, and that Jones had syphilis and smoked. Does it matter for our understanding of Jones's case how many of the other people who have syphilis also smoke, and thus, what the probability p is? No; that fact is surely irrelevant to our understanding of this *particular* case. The exact same reasoning holds if we substitute an unknown property H for smoking.

There is of course the question of how justified we are to believe that untreated syphilis caused paresis in the case of Jones. But our degree of justification is not at all related to the value of p , at least not if we are justified in believing that untreated syphilis is a necessary condition for paresis. If there are different possible causes of paresis, then the value of p may become relevant to our epistemic state: if Jones smoked and has syphilis, and smoking causes paresis with a 40 percent chance while syphilis causes it with a 1 percent chance (very low p), we are not justified in believing that the syphilis was the cause of the paresis. But this is clearly not what the size debate was about.

Let us now look at Jeffrey's and Salmon's example of the child born of two heterozygous brown-eyed parents. I already quoted Salmon as saying that the birth of a brown-eyed child "can presumably be explained statistically on the basis of the 0.75 probability of such an occurrence." On my theory, the explanation of the birth of a brown-eyed child would be something like

this (we are still assuming that the underlying processes are deterministic):

The father and the mother are heterozygous.

$\exists H$: person a is heterozygous $\wedge H(a) \implies$ the ‘blue allele’ is inherited.

$\neg(H(\text{father}) \wedge H(\text{mother}))$ (*rather than this not being the case*).

If and only if the child inherits two ‘blue alleles’, it has blue eyes; otherwise it has brown eyes.

The child has brown eyes (*rather than blue eyes*).

Here, H is a (perhaps bewilderingly complex) statement about causal factors determining how meiosis in the gametes works and which gametes finally fuse. We do not know H , but in so far as we are justified to believe that the process is deterministic, we are justified to believe that such an H exists. Notice that probabilities do not play a part in this explanation. It doesn’t matter that the probability of $\neg(H(\text{father}) \wedge H(\text{mother}))$ is 0.75; all that matters is that it is *true*.⁴

But of course, that means that we can as easily explain the birth of a blue-eyed child, and that we will understand that event just as well. The explanation is simply:

The father and the mother are heterozygous.

$\exists H_a$: person a is heterozygous $\wedge H_a \implies$ the ‘blue allele’ is inherited.

$H_{\text{father}} \wedge H_{\text{mother}}$ (*rather than this not being the case*).

If and only if the child inherits two ‘blue alleles’, it has blue eyes; otherwise it has brown eyes.

The child has blue eyes (*rather than brown eyes*).

The only difference is that the third premise has turned into its negation; apart from that, the explanations are identical. There is thus, on the determination theory, no serious difference between explaining that the child has

⁴We could also explain the brown eyes without using probabilities, by simply stating that the child has at least one brown allele, and that this is a sufficient condition for having brown eyes. That would be a good explanation as well, but less useful for the purposes of the current chapter.

blue eyes and explaining that the child has brown eyes. Thus, the determination theory can explain and vindicate the intuitions of Salmon, Jeffrey and the later Hempel.

But what about the intuitions of Mellor, Strevens and the earlier Hempel? Their basic idea is that if explanation X explains why A rather than A' , that very same explanation cannot possibly also explain why A' rather than A . And indeed, it seems to be entirely rational to claim that we cannot explain a contrast by something that does not make a difference.

This is a problem for the SI model (and other models, such as the statistical relevance model). The explanation of “brown rather than blue” and the explanation of “blue rather than brown” are, in the SI model, exactly the same (except for the value of p) – and this is problematic. But in the ND model, these explanations are different: one of the premises turns into its negation, and it is exactly this premise that makes the difference. Thus, Mellor could accept my account where he could not accept that of Salmon; we have saved his intuition.

The answer to Strevens’s claims is perhaps somewhat different. Strevens asks us not whether the same *explanation* could explain both that heat always flows from hot to cold objects and that heat always flows from cold to hot objects; he asks whether the same *theory* can do so. This does not seem to be problematic. If, contrary to our expectations, heat always flowed from cold to hot objects (and we were nevertheless justified to believe that statistical mechanics is true), we could explain this fact in something like the following way:

According to statistical mechanics, if the actual state of the Universe is in this very small part P of the Universe’s phase space, heat will always flow from cold to hot objects.

Statistical mechanics is true.⁵

The actual state of the Universe is in P .

Heat always flows from cold to hot objects.

And rather than this being a fanciful kind of explanation that nobody would ever accept, an explanation exactly analogous to this is frequently given to explain the second law of thermodynamics. See, for instance, Albert 2000 [2] on the Past Hypothesis. The idea here is that in order to explain the

⁵Or empirically adequate, for those so inclined, although it’s hard to see how it *could* be empirically adequate in the present example, even if it were true!

fact that we have time-symmetric laws of nature but a constantly increasing entropy, we must posit an extremely unlikely low-entropy state in the past – just as we would have to posit an extremely unlikely state to explain the opposite.⁶

My conclusion is that the ND model satisfies all the intuitions in the size debate.

Concerning the change debate, I claimed that its moral was that although counteracting causes can hardly be said to explain the event in question, we nevertheless often mention them in explanations, and believe that we understand what happened better when we do. Can we explain this feature of explanation using the ND model?

Let us give an explanation of the survival of a plant that has been sprayed with defoliant:

This plant was sprayed with defoliant.

$\exists H : (x \text{ is sprayed with defoliant} \wedge H(x) \implies x \text{ survives}) \wedge (x \text{ is sprayed with defoliant} \wedge \neg H(x) \implies \neg x \text{ survives}).$

$H(\text{this plant})$ (*rather than this not being the case*).

This plant survived (*rather than not surviving*).

(Normally, we might want to say something more about H , since we know that among other things it involves getting enough sunlight, water and minerals, not being eaten by goats, and so on. But let us ignore this for the moment.)

Cartwright's complaint that we cannot give an explanation merely by citing counteracting causes is immediately understandable given this construction. Giving only the first premise of this argument would not be a good explanation. Why not? The answer cannot be that giving only the

⁶Also note that the increased status of statistical mechanics after the work of Maxwell and Boltzmann – which Strevens cites as evidence for size elitism – can be explained in many ways. One explanation could be that the fact that Maxwell and Boltzmann could deduce high probabilities for processes actually seen increased the (epistemic) probability of their premises being true, which made people less reluctant to accept them. Strevens rejects this view, because he considers it absurd to claim that Maxwell and Boltzmann did not increase the explanatory power of statistical mechanics. So perhaps a better explanation of the higher status of statistical mechanics is that Maxwell and Boltzmann indeed increased the theory's explanatory power, not by deducing high probabilities (as the elitist would have it), but by developing the concepts we need to characterise the difference between those initial conditions that do, and those that do not lead to the observed behaviour. But this is not the place for a full treatment of this issue.

first premise is an abbreviation that leaves some of the explanation implicit, because we abbreviate explanations all the time. No, what is going wrong is that even an abbreviated explanation *must* always mention at least one of the *contrastive* premises; one of the things that made a difference. In this case, if we mention only the first premise, we do not mention the contrastive premise, which is the third premise; and thus we do not give a good explanation.

But why is giving a counteracting cause explanatory useful at all? Because it allows us to make the non-specific premise more specific, and the more specific it is, the better we understand what happened. The H that is a necessary and sufficient condition for a plant to survive if it has been sprayed with defoliant is much more specific than the G that is a necessary and sufficient condition for a plant to survive. We can see this by noticing that G could be written as “sprayed with defoliant and H , or not sprayed with defoliant and H' ”, for some H' .

What we have seen in this section, then, is that the determination theory and the ND model automatically save the intuitions of both sides in the size debate. We can also make excellent sense of the conflicting intuitions in the change debate if we assume that (a) even abbreviated explanations must mention one of the contrastive premises, and (b) making the non-specific premise in an explanation more specific increases understanding. Both these assumptions are plausible. I conclude that the determination theory can successfully account for the role of probability in explanations of deterministically produced events. (But we will return to this topic at length in chapter 8.)

7.4 Totally unspecific explanation

The ND model gives rise to an obvious problem, though. Assuming that E is a deterministically produced event, we could use the ND model to trivially answer the question why E happened by constructing what one could call the *totally unspecific deductive explanation* or TUDE:

$$\exists H : (H \implies E) \wedge (\neg H \implies \neg E).$$

H (rather than this not being the case).

E (rather than this not being the case).

In words, this would come down to saying that E happened because something determined it to happen. This is a very unsatisfying explanation,

because it gives us no understanding at all.

Different explanations of the same explanandum can give us different amounts of understanding; some may give a lot of understanding, while others give only a little. A TUDE is the limit case of an explanation which gives zero understanding. If we believe this means that a TUDE is not an explanation at all, we must add to the determination theory a condition number 6 stating that the explanation is not allowed to be a TUDE; alternatively, we can say that a TUDE *is* an explanation, albeit the worst possible explanation. This is a decision of vocabulary; nothing hinges on it. For reasons of simplicity, I will adopt the latter option.

The idea of explanations giving different amounts of understanding will be discussed in more detail in chapter 9.

7.5 Indeterminism: biting the bullet

Until now we have spoken of non-deterministic explanations of deterministically produced events. In such cases, we can apply the ND model. But of course there are – or at the very least there might be – cases of true indeterminism, cases where something simply *happens* without being determined by any factor. The least controversial examples of true indeterminism are quantum events such as the decay of a radioactive atom, and it is that example which I will use in this section.

Suppose we are looking at a ^{229}Np atom, an unstable isotope of Neptunium with a half-life of 4 minutes. After 6 minutes and 22 seconds, the atom decays. Can we explain that the atom decayed after precisely this interval? Can we explain that the atom decayed within an hour? That it decayed within one hundred years?

The determination theory implies that we cannot explain *any* of these facts. Since there is no factor that determined that the atom would decay at this rather than at another moment, there is nothing we can give as an explanation for the exact moment of the decay, or even for the decay taking place within three million years. This is hardly a novel position, and similar sentiments have been expressed by, among others, Mellor 1976 [80], who writes:

Sometimes, however, suitable causal explanation is not to be had.
An event may lack sufficient causes, as radioactive decay does.
(p. 235.)

and Van Fraassen 1980 [28], who says that the precise moment of radioactive decay is “the sort of fact that atomic physics leaves unexplained” (p. 108).

Does this mean we cannot explain anything about radioactive decay at all? Is quantum mechanics without explanatory power? No, of course not, for there are other explananda that we *can* explain using quantum mechanics. For instance:

Why was the probability that this ^{229}Np atom would decay within 8 minutes approximately 75% (rather than some other percentage)?

is exactly the kind of thing that quantum mechanical atomic theory sets out to explain and *does* explain. The following three facts are explained as well:

Why was the probability that this ^{229}Np atom would decay after approximately 6 minutes and 22 seconds greater than zero (rather than zero)?

Why was it possible (rather than impossible) for this ^{229}Np atom to decay after exactly 6 minutes and 22 seconds?

Why is it likely (rather than unlikely) that the emission rate of this sample of ^{229}Np will decrease by about 50% in 4 minutes?

The explananda in all of these explanatory requests are strictly determined by (can be deduced from) the laws of quantum mechanics (and appropriate other facts).

According to the determination theory, then, where true indeterminism reigns we can explain probabilities and possibilities (if we have an appropriate theory), but not the actual events themselves. This holds true even for statistical amalgams of individual events, such as:

Why did approximately half of these 10^{14} ^{229}Np atoms decay in four minutes?

for that is, after all, only one more actual event that, although very likely, was not necessary.

With respect to indeterministic explanation I simply bite the bullet and say that we do not understand why the atom decayed at a specific time, and we do not even understand why it decayed within one thousand times its half-life. When faced with an explanatory request concerning such an event, we ought to reject them. In practice, we generally explain something else – we explain why the event was possible, for instance, or why the event was probable. Once we have done that, we may in a loose sense say that we “understand something” about the event; but strictly speaking, we do

not understand the explanandum; we merely understand some facts modally relevant to the explanandum.

I thus agree with Jeffrey (1969) [53] that the probability we can deduce for the explanandum is not a measure of how well we understand it. Suppose an atom has a 0.99 probability of decaying within an hour, and a 0.01 probability of not doing so. Will we have a better understanding of what happened when it does decay then when it does not decay? No, says Jeffrey (and we have already seen in section 7.2 that Salmon and the later Hempel agree with him):

The strength of a statistical explanation . . . is not given by the degree of confirmation that the premises bestow on the conclusion in the corresponding Hempelian inductive inference. . . . The knowledge that the process was random answers the question, ‘Why?’ – the answer is, ‘By chance’. . . . the explanation is basically the same no matter what the outcome: it consists of a statement that the process was a stochastic one, following such-and-such a law. (One may gloss this statement by pointing out that the actual outcome had such-and-such a probability, given the law of the process; but this gloss is not the heart of the explanation.) ([53], p. 109.)

I only add to this that “By chance” is not an explanation at all, but rather an admission of explanatory defeat.⁷

7.6 Indeterminism: is it a bullet?

The claim that non-determined explananda cannot be explained has been vigorously rejected by some philosophers. For instance, Salmon 1984 [110] argues against “holding that events are explainable only to the extent that they are fully determined”. His argument is that quantum mechanics “must be admitted to provide genuine scientific explanations of a wide variety of phenomena” (p. 53). But of course the two claims are not contradictory, as my previous examples have shown. To explain why the probability of an ²²⁹Np atom decaying within 8 minutes is approximately 75%, is to provide a genuine scientific explanation of an interesting phenomenon. The determination theory does not belittle quantum mechanics.

⁷To say that something happened “by chance” is to say that it “*just* happened”, which is to say that it cannot be explained. It gives us more *information* – we now know that there was no deterministic process that led to the event – but it does not give us more *understanding*.

A related argument given by Salmon is that the advent of quantum mechanics has made it imperative to change our idea of what a good explanation is:

If we embrace indeterminism, we must adopt a suitable conception of explanation to go along with it. (Salmon 1998 [113], p. 42.)

But this is a non sequitur. A general theory of explanation must be able to answer the question “If E is the result of an indeterministic process, can it be explained?”; it is neither allowed to tell us that such processes exist, nor allowed to suppose that they do not. Whether or not we believe our world to be indeterministic is irrelevant to judging the success of a theory of explanation. This doesn’t mean that changes in science cannot *cause* changes in the philosophy of explanation; it is just hard to see how they could enter as *arguments* into the debates. There is no logical connection between determinism and the claim that only determined events can be explained. We can believe that the world is deterministic while also believing that indeterministic events could be explained; and we can believe that the world is indeterministic while believing that only determined events can be explained. A logical link could be forged only by something like the claim that everything that actually happens can be explained, i.e., something like the Principle of Sufficient Reason – which Salmon would not wish to defend. There is no reason to change the notion of explanation such that anything that is described by modern science can be explained.

Hitchcock 1999 [47] goes as far as to claim that the idea that indeterministically produced events can be explained “has become close to orthodoxy in the philosophy of science to believe that indeterministic explanation is possible” (p. 585). As evidence for his claim he produces an impressive sequence of articles that give “powerful arguments” that have helped us to repress our deterministic intuitions: Hempel 1965 [41]; Jeffrey 1969, [53]; Salmon 1971 [108], 1984 [110]; Railton 1978 [97], 1981 [98]; Lewis 1986a [67], 1986b [68]; and Humphreys 1989 [51].

I do not believe there is an indeterminist orthodoxy here. It would, for instance, be non-trivial to argue that Lewis 1986a (which does not talk about determinism at all) and Lewis 1986b (which defends only indeterministic *causation*) are meant to be combined into an account where explanation is possible in situations where indeterminism is true. Indeed, Lewis 1986b claims that there can be no explanation of why the particle decayed within an hour rather than not:

[W]e are right to explain chance events, yet we are right also to

deny that we can ever explain why a chance process yields one outcome rather than another. (p. 230.)

Railton 1981 does defend the possibility of explanation by probabilistic laws. Railton identifies an explanation with an “ideal explanatory text”, which is, more or less, a list of all the events and laws relevant to deducing the probability of a certain event. However, he writes:

[I]n 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 is simply 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. (p. 248.)

Thus, Railton claims that explanatory requests like “Why did the atom decay now, rather than at some other moment?” have no answer and are to be refused – which is exactly what I claim when I say that there is no explanation of the fact that the atom decayed now rather than at some other moment.

Lewis and Railton both believe that chance events can be explained, but that contrastive why-questions that ask why the event happened rather than not cannot be answered. Thus, there can be non-contrastive explanations of events that were not determined, but no contrastive explanations of why those events happened rather than not. Since I have argued earlier that all explanations are contrastive explanations, Lewis’s and Railton’s theses on indeterminism agree with mine. (We disagree about the existence of non-contrastive explanations, but this topic has been discussed in chapter 5.)

Hitchcock 1999 [47] sets out to show that Lewis and Railton are wrong in thinking that contrastive explanation is any different from non-contrastive explanation, and that in both cases non-determined events can be explained. His account, then, is particularly important for our discussion, and I wish to focus on his arguments for the rest of this section. However, one problem of coming to grips with it is that Hitchcock assumes that his readers already accept that indeterministic non-contrastive explanation is possible – but I do not accept this, since I do not believe that non-contrastive explanation exists. Nevertheless, he gives four counterexamples to my thesis that indeterministically produced (contrastive) events cannot be explained, and it will be instructive to look at these.

Hitchcock's first counterexample concerns the explanatory question "Why did Jones, rather than Smith, get paresis?". The answer is that Jones had untreated syphilis, and thus a small chance of getting paresis; while Smith did not have untreated syphilis, and thus no chance at all of getting paresis. Hitchcock believes this to be a satisfactory answer to the question. Hitchcock's second counterexample is a variant of the first, but more clearly indeterministic. The question is "Why did this photon, rather than that other one, go through the screen?", and the answer is that the first had (because of its polarisation) a small probability of going through the screen while the second had no chance at all.

Such examples were discussed in section 5.1, where I argued that rather than normal contrastive explananda, the why-questions in these cases involve a *contrast of parallels*. I also argued that a contrast of parallels probably involves giving a satisfactory answer to both of the contrasted parallels. In this case, that would involve giving an explanation of why the second photon did not go through the screen (rather than going through it), which is easy; but also giving an explanation of why the first photon did go through the screen (rather than not going through it), which is, on my account impossible. Hence, there is no good answer to this why-question.

I am perfectly happy with this conclusion, but in section 5.2 I offered an alternative for those who believe that Hitchcock's examples *are* good explanations. The alternative consists in saying that answering a contrast of parallels involves only giving a satisfactory answer to *one* of the contrasted parallels. This alternative can be used to reconcile the Hitchcockian intuition with the determination theory, since there *were* causes that *determined* that the second photon would not go through the screen.

Hitchcock's third and fourth counterexamples do not involve a contrast of parallels. In the third example, we explain why Lewis went to Monash rather than to Uppsala by pointing out that he had an invitation for Monash, but not an invitation for Uppsala, and that he never goes anywhere if he doesn't have an invitation. We accept this explanation, says Hitchcock, even though we do not know whether the invitation determined Lewis to go to Monash – the possibility of his staying in Princeton is left wide open. This does not invalidate the explanation, hence, determination is not necessary.

According to condition 2 of the determination theory, we are entitled to assume that only the elements in the contrast class are possibilities. The explanatory request "Why did Lewis go to Monash rather than to Uppsala?" presupposes that he would go to either, and we are allowed to use this as a premise in our explanation. Given the premises of our explanation, then, it was indeed determined that Lewis would not stay in Princeton; hence, the counterexample is not a counterexample at all. (Hitchcock might actually

agree with my criticism of this third example. A large part of his own article is devoted to developing the idea of explanatory presuppositions, including those of contrastive explanations, and everything he says there is compatible with what I say in this paragraph. What is more, in the example he discusses on page 604 he makes an argument almost identical to the one I make here.)

The fourth counterexample involves a photon polarised in the vertical direction. The probability of transmission of this photon is zero if the polariser is aligned horizontally, and increases to one as the polariser is turned towards the vertical. Now, suppose I believe that the polariser is aligned horizontally; the photon, however, passes through, and I am very surprised. “Why did the photon pass the polariser, rather than being absorbed?”, I ask.

I tell you that the polarizer was in fact extremely close to the vertical, making it vastly more likely that the photon would be transmitted than that it would be absorbed. Have I not answered your question? (Hitchcock 1999 [47], p. 592.)

The only way in which this example differs from the example of the decaying atom is that in this case, I believe that E is impossible, and E nevertheless happens. I am puzzled about the very *possibility* of the event – and of course this possibility can be explained, since the actual position of the polariser determines the event to be possible (and indeed likely) rather than impossible. Perhaps that is why the answer in the example may seem, at first glance, satisfactory: because it explains the possibility of the event. But the *explicit* question which I ask, why the photon passed rather than being absorbed, is not answered and cannot be answered. If this has been granted to us in the case of the decaying atom, it will also be granted here.

It seems then that the four counterexamples given by Hitchcock do not pose a problem to the determination theory. Of course, this does not prove that the determination theory is correct. All I have been trying to show in this section is that it is possible to hold that non-determined events cannot be explained; not that it is necessary to hold this thesis. This is, I think, enough: if the determination theory can handle all cases of explanation, it can be judged a success. There is no additional need to prove that no other theory can do so.

But let us spend a few moments looking at Hitchcock’s own alternative:

To provide a contrastive explanation is to provide information that is explanatorily relevant to the explanandum, given the presupposition that is expressed by the contrast. The notion of explanatory relevance given a presupposition can be modeled by,

although probably not reduced to, probabilistic relevance conditional upon the presupposition. (Hitchcock 1999 [47], p. 608.)

I accept everything Hitchcock says here as a *necessary* condition for explanation. If X rather than X' determines that E rather than E' , the occurrence of X must be probabilistically relevant to the occurrence of E given the presuppositions that $E \vee E'$ and $X \vee X'$. (The probability of E given X is 1, whereas this same probability given X' must be lower than 1.) I do not accept that Hitchcock's analysis is *sufficient*. In addition, we need a relation between X and E that allows for manipulation of E through X , and (more relevant in the present section) the probability of E given X (and the further elements of the explanans) must be 1. Until convincing examples are brought forward that show this final clause to be unnecessary, the determination theory can be accepted.

7.7 Conclusion

In this chapter, I have defended the determination theory against the charge that there are non-deterministic explanations. I will continue to do so in chapter 8, where I discuss Michael Strevens's claim that many determined events are best explained probabilistically even when a deterministic explanation is available.

In chapter 9 we will turn to several further elucidations of the determination theory. There we discuss whether all explanations are arguments, the role of laws and regularities in explanation, the plurality of explanations, and the distinction between explanation and understanding.

