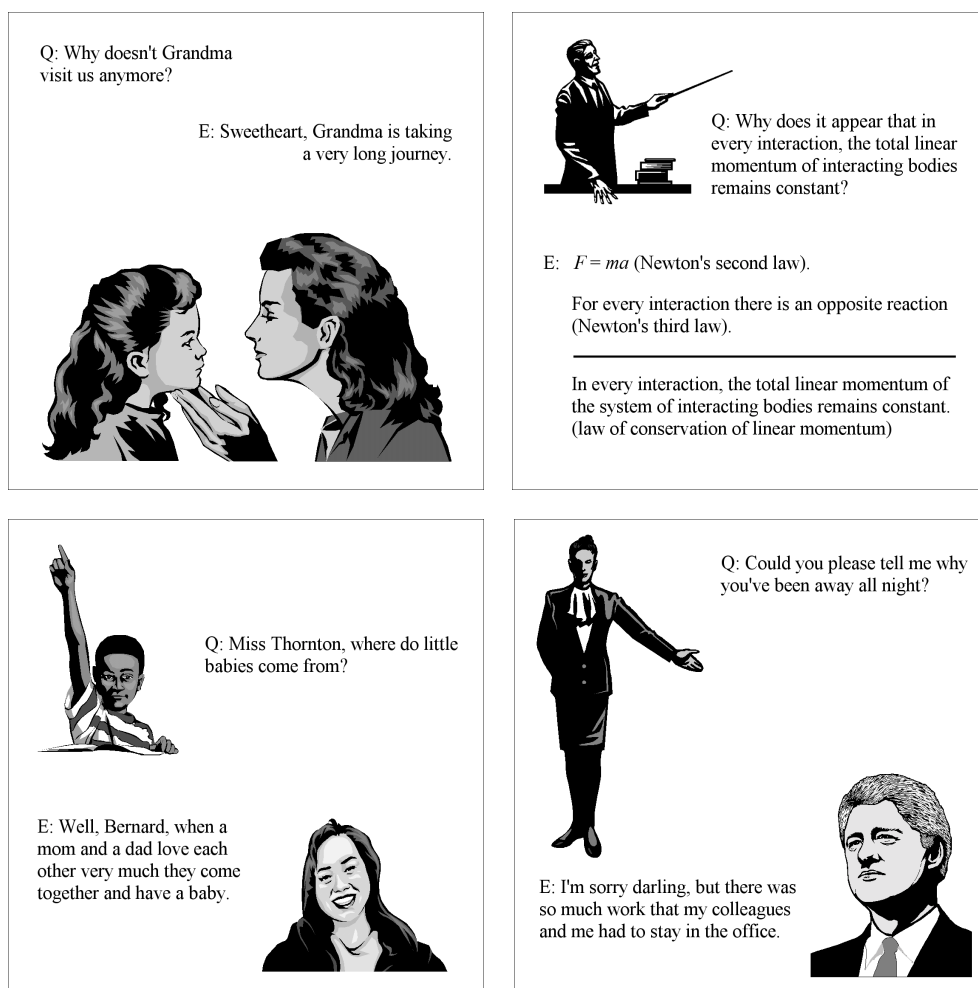# Explanations are like salted peanuts[*]

On why you can't cut the route toward further reduction

Daniel Cohnitz

1. The Pragmatics of Explanation

Take a look at these four situations:

Figure 1

All of these situations have certain features in common: in all of them an explanation is asked for, in all of them an explanation is given, and all these explanations are *literally false* (although in different ways).

That they are literally false can be *justified* in each of the four situations: there are moral, pedagogical, rational-egoistic, epistemic and other context-dependent reasons for us not to tell the literal truth sometimes. This is what I call *the pragmatic aspect of explanations*. Besides all the different adequacy

criteria that we have for an *ideal* scientific explanation, there is always a bundle of context-dependent pragmatic criteria that determine what counts as a good explanation in a certain situation.

Although all this is granted, there is nevertheless the idea of a normative, regulative model of scientific explanation, which is not context-dependent. We might be justified in making certain false claims in a situation in which we are asked for an explanation, but nevertheless we should not dream of calling these claims '*complete* scientific explanations'. One of these ideal explanatory schemas was described by Carl Gustav Hempel and further developed by Peter Railton. On both views there is a set of adequacy criteria that are meant to be necessary as well as sufficient conditions, given as an explication of the concept of scientific explanation.

A common feature of both accounts is that they demand truth for every sentence of the *explanans*. Railton is even more rigorous in this respect than Hempel, since the invocation of the truth-requirement leads him to dismiss Hempel's model for inductive explanations. In the next section (2) I shall give a brief outline of what we mean by an ideal scientific explanation and of the reasons that count in favor of such a regulative construction. The model (as I shall defend it) is in a certain way *mildly reductionistic*, i.e. it supports the following claims:

(I)   If some explanatory mechanism M can be reduced to some more fundamental mechanism M', any explanation E' referring to M' contains more information about the ideal explanatory text than any explanation E referring to M.[1]

(II)   If some explanation E' contains more information about the ideal explanatory text than some rival explanation E, E' has a higher explanatory value than E.[2]

Prima facie it does not seem very problematic to argue in favor of these claims. Depending on how we want to define an "explanatory mechanism", an "ideal explanatory text", and the notion of "explanatory value", both claims could simply turn out to be trivial tautologies. Nevertheless, Robert Batterman attacked the latter claim (which is indeed not vacuously true by definition, as we shall see) in various papers[3], trying to show by way of examples that explanations exist which cannot be surpassed by explanations revealing more of the ideal explanatory text. In section (3) I shall summarize Batterman's criticism and discuss his alternative model of scientific explanation. The arguments in section (4) are a defence of the ideal explanatory text-conception against Batterman's points.

2.   Ideal Scientific Explanations

The main idea behind the ideal explanatory text account is that all scientific explanations are deductive-nomological in structure. Another basic condition is the above mentioned truth requirement for all sentences contained in the *explanans*. The latter seems to be an overwhelmingly plausible constraint: the lies of figure 1 are of course all designed to *satisfy* the questioner (at least temporarily), nevertheless they do not fulfill the purpose the questioner expects them to (that's why we make such a big moral deal out of lying): they do not *answer* his question, they just make him stop asking. But complete scientific explanations should (at least) do the former, they should indeed answer our questions.

The idea that all scientific explanations have a deductive-nomological structure may need further explaining. Hempel's original account has to deal with many well-known counterexamples, all of which were designed either to show that Hempel's adequacy criteria exclude *bona fide* explanations, or to imply look-a-likes which are in no way explanatory.[4] So why stick to the deductive nomological idea and even expand it? Peter Railton thinks (on my view correctly) that Hempel's famous model needs only one additional requirement to be able to deal with all this criticism.

2.1   Hempel's Ideal

Hempel's explication of the concept of scientific explanation distinguishes between *deductive-nomological* (DN) explanations of general regularities, DN explanations of particular facts, *deductive-statistical* (DS) explanations of *general regularities*, and *inductive-statistical* (IS) explanations of particular facts. The first three are *adequate scientific explanations* in Hempel's mind iff (if and only if) they have the structure of a deductively valid and sound argument with the *explanandum* (regularity or fact to be explained) as the conclusion, and at least one (either deterministic or, in the case of a deductive-statistical explanation, probabilistic) law among its premises which is necessary to derive the conclusion from the set of premises (that form the so called *explanans*). The latter type, an *inductive-statistical* explanation, is an adequate scientific explanation iff it has the structure of an inductively valid (the premises support the conclusion with a probability $r = 1 - \varepsilon$ for some small number ε) and sound argument with the explanandum (the particular occurrence) as conclusion. Again it is demanded that it has at least one (statistical) law among its premises which is necessary to infer the conclusion from the set of premises, and, additionally, that this set contains the strongest information statistically relevant for the occurrence of the explanandum that is currently available relative to our background knowledge. The last requirement is also known as the *requirement of maximal specificity* (RMS).

It is obvious, even from this short summary, that IS-explanations play a rather unique role in Hempel's explication. They are not conceived as *de-*

ductive arguments and because of the requirement of maximal specificity they are even *nonmonotonic*: in a different epistemic situation, with new evidence about the statistically relevant factors for the explanandum to occur, new premises might undermine the support for the conclusion. Hempel's idea that predictions and explanations are structurally identical[5] leads him to demand a high probability ($r = 1 - \varepsilon$ for some small number ε) for the explanandum's occurrence relative to the explanans. The oddest consequence of this is surely that very unlikely chance events become *inexplicable in principle*, although the process that brought them about is in no way mysterious, but might be perfectly understood.[6]

2.2   Railton's Ideal

Railton's basic idea to address these oddities and counterexamples is that chance phenomena, even very unlikely ones, can be explained by subsuming them under irreducibly probabilistic laws in a deductive fashion. Hence the name of his account: *deductive-nomological model of probabilistic explanation* (DNP).

It should be emphasized that only *genuinely* probabilistic laws are considered as candidates for a DNP-explanans. In a strong interpretation this is already a consequence of Hempel's account[7], given RMS, the truth requirement, and Hempel's propensity interpretation of probabilities. But it is also intuitively evident: an event that merely appears to be a chance event *cannot* be explained statistically, for no probabilistic *laws* would govern it.[8]

This is still not enough to rule out the counterexamples directed against the DN-structure itself, counterexamples showing that adequate DN-"explanations" are not always explanatory. Instead of invoking some causal relevance criterion, Railton demands that the laws cited in the explanans are to be derived from our theoretical account of the mechanism(s) at work for the occurrence of the phenomenon to be explained. The reference to this mechanism(s) informs us about how the phenomenon came about, not why it was to be expected. Hempel stressed the latter point, whereas the first seems to be the one asked for.

These features (deductive structure, reference to the mechanism(s) at work) solve most if not all of Hempel's original problems with the IS-model: first, it is now possible to explain even unlikely chance events, second, RMS, and hence the relativity of an explanation to an epistemic situation vanish, and non-explanatory "explanations" are all filtered out. The latter was also a problem of the original DN-model, which is in Railton's account equally backed up by the mechanism(s) at work. Hence the schematic structure of Railton's DN-model of scientific explanation can be regimented like this:

(1)   Deductive-Nomological-model for the explanation of non-chance events

(1a)   A theoretical derivation of a deterministic law of the form (1b).
(1b) $\forall x\, \forall t\, [Fx,t \rightarrow Gx,t + \theta];\ \theta \geq 0$
(1c) Fe,t_0

(1d) $Ge,t_0 + \theta$

(1a) is a derivation of the deterministic law from our theoretical account of the mechanism(s) at work which guarantees the relevance and explanatoriness of the covering law. (1b) is the law itself, with the two place predicates *F* and *G* which are properties had by whatever the variable *x* or the constant *e* stands for at some time *t* (e.g., "*x* has the property *F* at *t*"), and the time-interval θ. (1c) states the initial conditions, and finally the explanandum (1d) is deductively derived from (1b) and (1c).

This differs from Hempel's structure only in (1a), the backup of the explaining law. Obviously more changes are to be made when it comes to the explanation of probabilistic events. Here the model looks like this:

(2)   Deductive-Nomological-Probabilistic-model for the explanation of chance events

(2a)   A theoretical derivation of a probabilistic law of the form (2b).
(2b) $\forall x\, \forall t\, [Fx,t \rightarrow \text{Prob}(Gx,t + \theta) = r];\ \theta \geq 0$
(2c) Fe,t_0

(2d) $\text{Prob}(Ge,t_0 + \theta) = r$
(2e) $(Ge,t_0 + \theta) / \neg(Ge,t_0 + \theta)$

Again (2a) embeds the probabilistic law into the mechanism at work, (2b) states the probabilistic law[9], and (2c) the initial conditions. (2d) ascribes a probability to a single case, derived from (2b) and (2c) via universal instantiation and *modus ponens*. The *parenthetic addendum* (2e) finally states how things turned out at $t_0 + \theta$. (2e) – the explanandum – is neither deductively nor inductively inferred from the explanans. But since it is a true chance event, it should not be derivable at all, by any set of initial conditions, and any empirical law.

We knew (2e) when we asked for an explanation, thus it doesn't carry any new information and is dispensable. Nevertheless, if we want to connect our explanation (2) with the explanation of any event caused by some instantiation of (2e) nearby, we will need the parenthetic addendum to connect the explanations in such a case. To model probabilistic explanations in this way means giving up their argument structure which was a special feature of Hempel's account. The explanandum is no consequence of the explanans-

premises anymore. At the same time this means losing the structural identity of explanations and predictions. But both consequences of Railton's account do not harm in any way the *explanatory virtues* of explanations so conceived.

The structures (1) and (2) are both to be understood as *ideals*. They provide the skeletal form of what we call *ideal explanatory texts*. An ideal explanatory text fully spelled out for some causal process would be

> an inter-connected series of law-based accounts of all the nodes and links in the causal network culminating in the explanandum, complete with a fully detailed description of the causal mechanisms involved and theoretical derivations of all the covering laws involved. This full-blown causal account would extend, via various relations of reduction and supervenience, to all levels of analysis, i.e., the ideal text would be closed under relations of causal dependence, reduction, and supervenience.[10]

Of course, such an ideal explanatory text can never be put down on paper. The text might be infinite (if time is continuous, even non-denumerably infinite) and hence not anything anyone actually demands from any particular scientific explanation. What is in fact demanded from a scientific explanation, is that it reveals parts of the ideal explanatory text, that it *carries information* about the ideal explanatory text. We call this *explanatory information*. Again the idea behind this notion is fairly straightforward:

Consider some ideal explanatory text for the explanation of fact *p*. Now consider any statement *S* that, were we ignorant about this text, but conversant in the language and the concepts employed in it and in *S*, *S* would reduce our uncertainty about the content or features of the explanatory text. If we model this reduced uncertainty by sets of eliminated possibilities, *S* provides *explanatory information concerning why p* to the extend *S* eliminates possibilities concerning the explanatory text.[11]

With the help of the latter notion we are now able to talk of *degrees of explanatoriness*. This allows us to identify the explanatory value of all *bona fide* explanations which are not 'complete'. Putting the pragmatic and circumstantial aspects of particular explanations in their right places, it seems that this account can cope with all tentative counterexamples perfectly. Well, Robert Batterman doesn't think so.

4. Batterman's Challenge

Consider the *chaos game*:

> [Mark off three] vertices of a triangle on a piece of paper. Label them 1, 2, and 3. Choose one point as your starting place, and start tossing a 3-sided die. [...] Suppose that your starting point is 1 and the first roll lands on the side marked 3. You must put a

point midway between your starting point 1 and the point marked 3. On the next roll, put a point halfway between this new point and the point assigned to the number rolled, and so on.[12]

After playing this game for a sufficiently long time, the following fractal pattern occurs, produced by apparently random placements of dots:
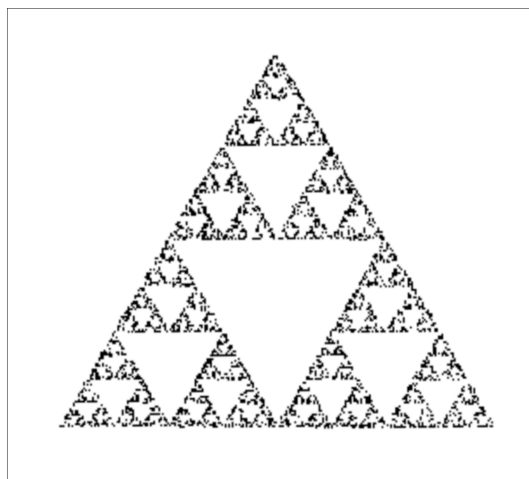

Figure 2. Sierpinski Triangle[13]

It appears that patterns resulting from different plays of the chaos game all instantiate the same fractal dimension. How do these patterns emerge? Why do they all have the same fractal dimension? According to Batterman, the answers to these questions cannot be provided by any ideal explanatory text.

Before I enter into the discussion of why Batterman thinks that phenomena like the patterns of the chaos game cannot be explained by some ideal explanatory text, let me add two methodological remarks:

1) The reader might be concerned about the fact that the chaos game seems more like a theoretical math problem than a counterexample taken from some empirical science. To be fair, Batterman takes the chaos game as a mere illustration of his point and even emphasizes the fact that it is "unphysical".[14] It must be noted however, that there is a way to treat the chaos game as a physical problem, i.e. the way Batterman treats the problem. Since this example is indeed illustrative, I will use it as well. The fact that it is in any sense "unphysical" will not harm the point I wish to make.[15]

2) Batterman wants to show that there are explanatory why-questions that *cannot* be answered by an ideal explanatory text. Thus it is not sufficient to show that there are explanations other than that which are *shorter*. As a matter of fact the ideal explanatory text-conception *emphasizes* that all actual explanations are shorter than the ideal explanatory text.

It indeed appears that Batterman found an explanandum in the chaos game which is not explicable by an ideal explanatory text. The question

(Q)   Why is it that, in general, sequences of dice rolls normally produce patterns like these (e.g. with such and such fractal dimension)?

seems to be an interesting "explanatory why-question", something we should have an answer for. But, quite remarkably, the correct answer is allegedly not the ideal explanatory text, the answer cannot even be formulated as an ideal explanatory text. How can we account for this?

The single results of throwing a die are usually not considered to be results of a irreducibly probabilistic process. In fact we often hear that, given complete information about all initial and boundary conditions and detailed knowledge of the throwing mechanism, we can accurately predict the result of each dice throw with certainty. Unlike irreducibly probabilistic processes, we assume hidden parameters here. Given a propensity interpretation of single case probabilities, to ascribe the probability 1/6 to the result that the die will land on "2" the next throw (again, with a normal die), is literally false: since the process is assumed to be fully determined by the initial and boundary conditions and the laws of classical mechanics, the probability that the die lands on whatever is only either 1 or 0.

Based on these assumptions, Batterman tries to show the insufficiency of the ideal explanatory text-model for explaining the observed regular patterns. Since the underlying process of a dice roll is assumed to be essentially deterministic, the ideal explanatory text cannot be probabilistic, and thus must be DN (compare schema (1)). But what should a DN-text that could account for the general pattern look like? Each dot in a Sierpinski triangle is the result of a single dice throw. Hence an explanation of any certain dot should consist of a deductive nomological explanation of the result of the roll, plus the rules of the game. By connecting all DN-texts together for each single roll and dot, we will get one large explanatory text that explains the occurrence of a single token of this pattern, the result of one play of the game. Is this the explanans we were looking for? Of course not. This would explain the coming about of a *token* of this pattern of such and such fractal dimension, not the general realizations of a *type* we were interested in. Indeed this latter question, as expressed by Q, is informative: not all *conceivable* sequences of dice throws would result in this pattern.[16] It seems that the only way to account for this phenomenon with an ideal explanatory text would be to ascribe a probability to each sequence of dice throws (each play of the game) to yield our pattern. But what underlying causal mechanism postulated would make such a theory derivable? Unfortunately there is no known theory that could satisfy this demand.

This is, essentially, Batterman's counterexample against the ideal explanatory text-conception. He then goes on to present how we in fact ex-

plain the general occurrence of these patterns using the Strong Law of Large Numbers (SLLN). First it is noted that the sequences of dice throws are Bernoulli sequences, i.e. every probability of a given outcome on a given trial is independent of the probability of any other trial in a sequence of throws. This makes the laws of large numbers applicable:

> Using the SLLN, it is possible to provide an apparently I-S-like explanation for the generic appearance of the triangle pattern by demonstrating that such patterns are highly probable in an ensemble of sequences of the kind described. In fact, one argues that the probability of realizing an instance of the pattern is one.[17]

Note that this "probability one" assertion in this "I-S-like" explanation has a measure theoretic meaning. It is not synonymous with "with certainty" (since we can conceive certain orderly sequences that would not yield the pattern). Such probabilities cannot be part of ideal explanatory texts. Probabilities can enter into an ideal explanatory text only if the phenomenon in question is irreducibly probabilistic (which is, by assumption, not the case).

Batterman goes on to show that physics is full of examples of structurally identical explanations. Two features are common ingredients of all these cases:

(i)   An assessment of the probabilities that the initial and boundary conditions are distributed in a certain way (analogous to our statement that a sequence of dice throws has the Bernoulli property).

(ii)   A proof that under iteration or toward a certain time limit a process starting with the probabilities for the distribution of initial and boundary conditions mentioned in (i) results in a certain macrophenomenon (analogous to the application of SLLN in the Sierpinski-example).

On the ideal explanatory text account one would have to question the status of the probability assertions in (i). To clarify whether and where they occur in the text, we have to see whether they are irreducible. Batterman's examples are all within the classical domain, hence they are clearly not irreducible. So we would have to replace them with DN-accounts for the occurrence of the particular initial and boundary conditions. After this replacement (ii) would not apply anymore (the probabilities the law would depart from are gone). Instead we would replace (ii) with the covering deterministic law. Again we arrive at (complete) explanations of singular occurrences, but would not yet arrive at an account for the properties shared by all these occurrences. If we wanted to regiment Batterman's alternative model, we would get something like this:

(3)   Statistical-Deterministic-model for the explanation of non-chance events

(3a)   A demonstration, through an analysis of the system's (*e*) lawlike dynamical instability, that *e* possesses certain strong statistical properties *SP* which results in (3b).

(3b)   *SPe,t₀*

$$(3c) \quad \forall x \, \forall t \, [SPx,t \rightarrow \text{Prob}(Gx,t + \theta) = 1]; \; \theta \geq 0$$

$$[r = 1]$$

(3d)   *Ge,t₀ + θ*

Again, (3b) corresponds to the Bernoulli property of the sequences of dice throws, (3c) to SLLN, which allows us to infer (inductively) with "probability one" that our system will exhibit the crucial pattern (as stated in (3d)). At least this is how I would interpret Batterman's own explication of his model:

> [The theoretical derivation from underlying theory as demanded by the DNP model] corresponds to the demonstration [(3a)], through an analysis of the system's lawlike dynamical instability, that the system possesses strong statistical properties [(3b)]. […] The second important feature […] is the derivation of a "probability one" assertion from a statement to the effect that the system has such a statistical property, together with a mathematical proposition; namely a measure-theoretic limit theorem such as the SLLN [(3c)]. Third, there is something at least intuitively satisfying about the I-S model in its claim that the explanans should provide us with evidence from which the explanandum [(3d)] can be inferred with high probability or with "practical certainty" [(inductive support (*r* = 1))].[…] Thus, the final component of the new model of statistical explanation is the inductive argument […][18].

Surprisingly the explanandum (3d) isn't not equivalent to Q! Q asked for an explanation of a regularity. (3d) is a sentence describing the fact that the system *e* has the property *G* (such and such fractal dimension) at $t_0 + \theta$. The inductive step is not part of the explanandum expressed by (3d), but the general occurence of our pattern was part of the explanandum expressed by Q. One is a little surprised to learn that Batterman's own model is just another (though shorter) explanation of one token of the pattern, an explanation that could admittedly be provided by an ideal explanatory text as well.

To be fair, throwing the inductive argument-idea overboard provides us with the explanation of Q. If we want an explanation for the *general* occurrence of our pattern, we might be able to show that *all* systems *e* of type *E* (plays of the game, in our example) exhibit the crucial statistical properties and hence we would arrive at a lawlike generalization similar to (3c), but quantifying over systems of type *E*.

It is absolutely true that no ideal explanatory text consisting of connected DN-explanations for the result of each single dice throw and dot could ever result in such a lawlike generalization. Luckily it doesn't have to. The versions of ideal explanatory texts that are discussed by Batterman are designed to account for particular facts, not for lawlike generalizations. Neither Railton's deductive-nomological or deductive-nomological-probabilistic model, nor Hempel's inductive-statistical model are applicable, since they are simply not intended to cope with statistical generalizations as explananda. The model that would be adequate to Q is a *deductive-statistical* explanation, for only here do we find an explanation of statistical regularities. It is not very amazing that explicatory models for scientific explanations of particular facts do not cover scientific explanations of regularities.

Could we provide a deductive-statistical explanation for Q in the sense of an ideal explanatory text? Maybe Batterman has shown that we cannot. Let us again consider his explanation (the way we have just redescribed it):

(4)   Statistical-Deterministic-model for the explanation of regularities

(4a)   A demonstration, through an analysis of the lawlike dynamical instability of systems e of type E, that E possesses certain strong statistical properties SP.

(4b) $\forall x \, \forall t \, [SPx,t \rightarrow \text{Prob}(Gx,t + \theta) = 1]; \theta \geq 0$

(4d) $\forall e$ of $E \, \forall t \, [SPe,t \rightarrow \text{Prob}(Ge,t + \theta) = 1]; \theta \geq 0$

(4a) certainly contains non-irreducible probabilities, since all considered cases stem from the classical domain. Hence, on the ideal explanatory text account we would have to replace them with laws, which are non-statistical, but allow for the derivation of (4d). How should this be possible? It isn't possible, but again, luckily it doesn't have to. Let's have another look at the explanandum: (4d) contains probability assertions that are not irreducibly probabilistic. Batterman would not claim that they are, since his point is (as reconstructed here) to emphasize that these reducible probabilities are derivable from regularities which equally contain reducible probabilities, and whose reduction into a DN-text would not be explanatory anymore.

From the explanatory text point of view, (4d) is not a true statistical law, but it is either false, or it is a non-statistical law, or it is not a law at all, simply *because* its probability assertions are non-irreducible. Take a look at all three interpretations of (4d):

(a)   The explanandum is false.
If (4d) is false, it is not surprising that we cannot derive it from true premises. This cannot be a challenge to our account, since no one would expect a theory of scientific explanation to provide explanations for falsities. No

ideal explanatory text can account for a false explanandum, but a false explanandum should not be explicable by any model of scientific explanation.

(b)   The explanandum is a non-probabilistic law.
Consider (4b) to be a shorthand for a more complex non-statistical law. This would be possible, but at the same time it would mean that its more complex relative is derivable from other (complex) non-statistical laws. Fully understanding the underlying mechanism would provide us with both. Again, this can't be considered a challenge.

(c)   The explanandum is not a law
Consider (4b) to be a mere descriptive report of observed relative frequencies. In this case we already know that a conjunction of DN-texts is sufficient to explain this observed relative frequency. No challenge here either.

Whatever you choose, we can account for it in the ideal explanatory text-conception. I think Batterman has not succeeded in finding explananda which must be judged inexplicable from our point of view, but explicable from the view of some other model of scientific explanation.

5. Explanations are like salted peanuts ...

As a conclusion of the previous argument, let me briefly answer three questions (for the record) that might be left open:

(1)   Do I say that all explanations which refer to some ergodic theorem or that involve Bernoulli properties and laws like SLLN are non-explanatory?

Surely not. On the contrary, I think that we can identify their explanatoriness very well[19]. Railton himself shows how ergodic theory can carry explanatory information[20]:

> For example, various proofs in ergodic theory and related results show that if a gas is in an initial condition that obeys a relatively few constraints, it will, over infinite time, spend most of its time at or near equilibrium. This illuminates a *modal* feature of the causal process involved and therefore a modal feature of the relevant ideal explanatory text [...].[21]

Batterman doubts that this strategy works, since the "modal feature" Railton addresses would not be a feature of the causal process, but a feature of the distribution of initial and boundary conditions.[22] It might be that some limit theorems exhibit modal features of the causal process, i.e. robustness against

a certain range of variation among the initial and boundary conditions. Some results in game theory could be interpreted in this way.[23] Or we may find the regularity in the distribution of the initial and boundary conditions themselves. Either way: ergodic theory might be a helpful *explanatory* shorthand for things that surely would take a while to be fully spelled out in an ideal explanatory text. But again: this text is an *ideal*. No one expects any actual explanation to be *complete* in this sense. What I had to show was simply that the ideal explanatory text wouldn't leave out something important some other explanation could include.

(2)   Do I say that all these statistical-deterministic explanations are literally false?

In a way I do. Just as some of the explanations in figure 1, these statistical-deterministic explanations could be misleading if taken literally. Knowing that the ascribed probabilities are meant in a measure theoretic sense, or represent statistical distributions of initial conditions only, they are nevertheless perfectly okay: they tell us in what direction we have to look for a "deeper" explanation.

(3)   Has Batterman failed to present a counterexample because he attacks a tautological, interdependent net of definitions?

No, he indeed attacks a substantial claim. My presentation of the ideal explanatory text-conception shows that the "mildly reductionistic" claim (I) is indeed a mere definition. Claim (II) on the other hand is contentful: "explanatory value" is an intuitive notion that we used to compare the two conceptions of scientific explanation. (II) would have been falsified, if Batterman had shown that there are explanations of events the ideal explanatory text couldn't account for. This didn't happen though. What (II) nevertheless implies (and what Batterman failed to undermine) is that you cannot cut the route toward further reduction if you know that the underlying mechanism is different from the one you refer to in your explanans. It is always legitimate and explanatory to keep asking why-questions, leading in the same direction, toward more fundamental levels (if there are any). It is in this sense that explanations are like salted peanuts. Getting one doesn't make you stop asking for one; usually just the reverse is true.

Notes

Marcus Rossberg, Gerhard Schurz, Elliot Sober, Barry Smith, Rein Vihalemm, the members of the philosophy of science section at the GAP conference, the members of the philosophy of science section at the XX Baltic Conference on the History of Science, and the members of the Estonian Association for Analytic Philosophy for helpful comments on earlier drafts of this paper.

[1] Note that this kind of mild reductionism doesn't include the claim that for every explanatory mechanism *M* there is some microphysical mechanism *M'* such that *M'* explains *M*, nor such that *M* is reducible to *M'* or anything like that. It is a mere *epistemologically* motivated principle that implies nothing about the *ontology* of our world.

[2] *E* might of course have a couple of other merits than explanatory ones which make it more appropriate in a certain situation. But this is pragmatics again.

[3] See Batterman 2000a, Batterman 2000, Batterman 1998, Batterman 1992. I'll concentrate on the argument presented in Batterman 1992.

[4] For the problems Hempel's account has to face, see Salmon 1990 and Schurz 1995/1996.

[5] Plus his above mentioned idea that all explanations are basically either deductively or inductively valid arguments with the explanans as conclusion.

[6] Railton 1978, 212: "After all, why should it be explicable that a genuinely random wheel of fortune with 99 red stops and 1 black stop came to halt on red, but inexplicable that it halted on black? Worse, on Hempel's view, halting at any *particular* stop would be inexplicable, even though the wheel must halt at some particular stop in order to yield the explicable outcome *red*."

[7] Although it took him a while to see it. Railton 1978, FN2.

[8] Railton himself distinguishes between deterministic and indeterministic systems in the following way: "Let us say rather loosely that a system is deterministic if, for any one instant, its state is physically compatible with only one (not necessarily different) state at each other instant. A system is indeterministic otherwise, but lawfully so if a complete description of its state at some one instant plus all true laws together entail a distribution of probabilities over possible states at later times." Railton 1978, FN1.

[9] Read "Prob(*Gx,t* + θ) = *r*" as "the probability that *x* has the property *G* at *t* + θ equals *r*".

[10] Railton 1981, 247.

[11] Since 1981 there has been progress towards a semantic theory of information. I won't go into details since I think that the intuitive grasp of the idea should be sufficient for the point I wish to make. For a detailed analysis I recommend Perry/Israel 1990.

[12] Batterman 1992, 325.

[13] Taken from Batterman 1992, 326.

[14] Batterman 1992, 325.

[15] To be precise, it won't harm my point to any higher degree than Batterman's. If it isn't a contentful explanation, but a mere mathematical example, the coming about of the sequences considered doesn't make a differ-

ence. But in this case the example doesn't count as a counterexample and is worthless as an illustration for Batterman's point, too.

[16] If the die always landed on "2", or if we had repeating sequences of "1", "2", "3", "1", "2", "3", ...

[17] Batterman 1992, 333.

[18] Batterman 1992, 343-344.

[19] One has deeply misunderstood the explanatory text-account if one thinks that Railton claimed that the statistical regularities of statistical mechanics wouldn't be *explanatory*.

[20] I do not wish to enter the discussion whether ergodicity and ergodic theory have any significance for explaining why equilibrium statistical mechanics works. Compare Batterman 1998, Sklar 1973, Earman/Rédei 1996.

[21] Railton 1981, 251.

[22] Batterman 1992, 345.

[23] Batterman himself addresses this point in D'Arms/Batterman/Górny 1999.

References

Batterman, R. (2000a): Multiple Realizability and Universality, in: *British Journal for the Philosophy of Science* 51, pp. 115-145.

Batterman, R. (2000): A 'Modern' (= Victorian?) Attitude Towards Scientific Understanding, in: *The Monist* 83, pp. 228-257.

Batterman, R. (1998): Why Equilibrium Statistical Mechanics Works: Universality and the Renormalization Group, in: *Philosophy of Science* 65, pp. 183-208.

Batterman, R. (1992): Explanatory Instability, in: *Noûs* 26, pp. 325-348.

D'Arms, J./Batterman, R./Górny, K. (1998): Game Theoretic Explanations and the Evolution of Justice, in: *Philosophy of Science* 65, pp. 76-102.

Earman, J./Rédei, M. (1996): Why Ergodic Theory Does Not Explain the Success of Equilibrium Statistical Mechanics, in: *British Journal for the Philosophy of Science* 47, pp. 63-78.

Perry, J./Israel, D (1990): What is Information?, in: P. Hanson (ed.): *Information, Language and Cognition*, Vancouver: University of British Columbia Press, pp. 1-19.

Railton, P. (1978): A Deductive-Nomological Model of Probabilistic Explanation, in: *Philosophy of Science* 45, pp. 206-226.

Railton, P. (1981): Probability, Explanation, and Information, in: *Synthese* 48, pp. 23-256.

Salmon, W. (1990): *Four Decades of Scientific Explanation*, Minneapolis: University of Minnesota Press.

Schurz, G. (1995/1996): Scientific Explanation: A Critical Survey, in: *Foundation of Science* I/3, pp. 429-465.

Schurz, G. (1993): Scientific Explanation: A Critical Survey, in: G. Schurz/A. Hieke (eds.): *IPS-Preprints* 1.

Sklar, L. (1973): Statistical Explanation and Ergodic Theory, in: *Philosophy of Science* 40, pp. 194-212.