

Indeterminism and the contrastive theory of explanation

Petri Ylikoski

Developments in physics during the last century have forced us to give up the idea that the world is fully deterministic. This has led philosophers to consider the concept of indeterministic causation. It is no longer plausible to think that all probabilistic claims in science are only reflections of our ignorance. There are at least some contexts in which the probabilities can be plausibly taken as reflecting the truly indeterministic nature of the processes under consideration. In this paper, I will take this situation for granted and proceed to consider its implications for theory of explanation.

The theory of explanation I will be considering is the contrastive counterfactual model of explanation developed in Ylikoski 2001. Its basic idea is that the *explanandum* is always best seen as contrastive. An explanation is an answer to a question that asks why *f* occurred instead of *c*. It tells us what made the difference between *f* and *c*. The central benefit of the contrastive approach is that it allows one to be specific about one's *explanandum*. It can be used in evaluating what an explanation *can* explain. In this paper I will argue that indeterminism does not create any fatal problems for the contrastive approach, and that in fact the contrastive ideas are indispensable in making sense of indeterministic *explananda*.

1. Explaining outcomes of indeterministic processes

Wesley Salmon has argued that in some indeterministic contexts the *explanandum* is not contrastive. He argues on this basis that the contrastive theory does not apply to an important class of *explananda*. (Salmon 1998: 328-329.) He notes that contrastive explanation rests on the following principle:

If *a* explains *f* in one case, and *c* and *f* are incompatible, then *a* cannot explain *c* in another (similar) case.

Salmon calls this *the Leibniz principle* (after Stegmüller). David Ruben calls this same principle *Plato's second principle* (Ruben 1990: 64-66). Since the above formulation is in non-contrastive form, let us reformulate it as follows:

If *a* explains *f*[*c*] in one case, then *a* cannot explain *c* [*f*] in another (similar) case.

According to Salmon, this principle does not hold in the indeterministic cases that are familiar from modern physics (Salmon 1998: 154-158, 326-329). To understand Salmon's reasoning, let us consider an example. The example is not from a textbook of modern physics, but it allows us to see the problem clearly.

Example 1: A chance set-up

We have a coin-tossing device, with two coins and a switch. When the switch is down, a biased coin is tossed, so that the chance of heads will be .35 and the chance of tails

will be .65. When the switch is up, another biased coin is tossed, so that the chance of heads will be .65, and the chance of tails will be .35. We will suppose for the sake of simplicity that tossing a coin with this set-up always includes irreducibly random elements. We want two cases to be explained. In the case A, the switch is turned down and the result is heads. In the case B the switch is turned up and the result of a toss is also heads. Let the *explanandum* be $f[c]$ in the case A, and $f^*[c^*]$ in the case B.

According to Salmon, we cannot explain why the result of the throw in cases A and B was heads rather than tails. However, he claims that we can explain why the result was heads. So, we cannot explain why $f[c]$ or $f^*[c^*]$, but we can explain the plain *explananda* f and f^* . (For a similar account see Lewis 1986: 230-231.) This observation convinces Salmon that 1) the Leibniz principle is wrong in indeterministic contexts, 2) that not all *explananda* are contrastive, and 3) that we can also explain events with low probability.

My concern here is only with theses 1) and 2). According to Salmon, subscribing to the Leibniz principle and to the idea that all explanation is contrastive springs from an attachment to a deterministic worldview. If we give up this worldview, as modern physics suggests we should do, we should also give up the Leibniz principle and the contrastive account of explanation which rests on it. (Salmon 1998: 327-328.) I do not share Salmon's assessment of the situation. Remember that we are now speaking about explanation, not causation. It is quite uncontroversial that modern physics describes the world as including irreducibly indeterministic processes. However, I claim that the implications of this fact for the theory of explanation are not that obvious. Salmon needs to show that when we give up the deterministic picture of the world, we also need to give up these two fundamental ideas about explanation.

Salmon argues for his position by trying to show that we can do without the Leibniz principle. He claims that the main reason for keeping the Leibniz principle sacrosanct is that it rules out some pseudo-explanations. He gives three examples of pseudo-explanations: 1) the explanation of the power of opium to produce sleep by its dormitive virtue; 2) the explanation of all happenings by the will of God; and 3) the explanation of behavior by a psychoanalytic theory that is compatible with all possible behavior. Salmon argues that we don't need the Leibniz principle to show that these explanations are not scientific explanations. According to him, the dormitive virtue theory is too *ad hoc* to be acceptable, the will of God is a scientifically unacceptable explanation because it refers to a supernatural agency, and that psychoanalytic theories which are compatible with all possible behavior cannot be scientifically confirmed. According to him, it suffices that we require that scientific explanations appeal to *bona fide* scientific laws and theories. (Salmon 1998: 327.)

Salmon's argument is interesting. He claims that the Leibniz principle is only used to rule out the kinds of explanations he mentions, and that he can show them illegitimate without it. This suggests that the Leibniz principle is only an *ad hoc* device to rule out pseudo-explanation, not a deep principle of explanation at all. Let us see how his argument fares, before considering

whether the Leibniz principle has also some other uses.

As I see the situation, we don't need the Leibniz principle to rule out the dormitive virtue theory. However, Salmon's diagnosis of the situation is not accurate either. The explanation is not a pseudo-explanation because it is *ad hoc*. It is a pseudo-explanation because it is non-explanatory, and it is non-explanatory because it is uninformative. In contrast, some *ad hoc* explanations are both informative and explanatory. The trouble with the dormitive virtue explanation is that it simply describes the cause as the kind of cause that causes the effect we are interested in. This explanation does not require any empirical research, and it cannot fail. The *explanans* is basically deduced from the *explanandum*, and we do not want explanation to be that easy.¹

What about God's will as an explanatory factor? Here I think Salmon's diagnosis is *ad hoc*. It is not enough to claim that the reference to a supernatural agency is not acceptable in science: there should also be some kind of a justification for this claim. We would like to have some principled reason why the God's will is not an acceptable explanatory resource in science. Salmon owes us a true diagnosis of this case that does not presuppose the Leibniz principle. Before that, I think we should stick with the standard justification, which is the following: 'It's God's will' is too easy an explanation. It can explain everything equally well, and it is compatible with everything. Consequently it does not tell us why *f* happened instead of *c*. Because of this, its explanatory power is only apparent. The competing (scientific) theories do not claim to explain everything, but they at least sometimes provide us reasons why *f* had to occur instead of *c*. We are here basically resting on the Leibniz principle. So, when we appeal to the principle Salmon wants to dispense with, we can make sense of the situation.

I think that a similar solution is applicable to psychoanalytic theories that are compatible with all possible behaviors. (I don't know if there are such psychoanalytic theories, but let us suppose that there are.) If a theory really cannot explain any $f[c]$, it cannot explain anything. The point of scientific theories is to track dependencies in the world, and if a theory does not find any, its explanatory power is zero. We do not need considerations concerning the confirmability of the theory to establish this, as Salmon thinks. And if we do so, we again confront the charge that the reply is *ad hoc* in nature. We want an account which shows why the theories that are too easy to confirm are not good explanations. And for this purpose we need more fundamental principles, like the Leibniz principle.

Of course, Salmon might just have been too hasty in his analysis of these examples. However, there is a more general reason why his approach to these pseudo-explanations does not work. He thinks that the criterion for ruling out these pseudo-explanations is their unscientific nature. But if the argument rests on this criterion, one needs a criterion of demarcation between science

¹ There are some contexts where a reference to dormitive power can be informative and explanatory. If we do not know that opium has such a power and we want know why a person, who has had a dose of opium, fell asleep, the reference to opium's dormitive power is appropriate. However, a reference to a dormitive virtue as an answer to question 'why does opium cause sleep?' is completely unexplanatory.

and non-science. Salmon has not provided us any indication as to how this perennial problem can be solved without the Leibniz principle.

Furthermore, it is unclear why we should raise the issue of demarcation at all. An appeal to *virtus dormitiva* is a bad explanation outside science as well. We should be analyzing what makes explanations good in general. There are good explanations that are not scientific, and there are bad scientific explanations. It would be inconceivable, at least to me, if scientific and everyday explanations did not share some basic ideas about good and bad explanations. After all, science has its origin in common sense cognition. There should be some kind of continuity. Of course, it now includes much that is apparently incompatible with common sense, but I seriously doubt that this incompatibility also extends to the general principles of explanatoriness.

The above considerations suggest that Salmon is not successful in dispensing with the Leibniz principle. It is needed to explain why some pseudo-explanations are not explanatory. I think the real motivation for accepting the principle is that it reflects our conviction that an explanation should trace objective relations of dependency (see Ylikoski 2001). And since Salmon's ontic conception of explanation is based on this same conviction, it seems clear that he should also stick to something like the Leibniz principle.

Does Salmon have any other arguments for dispensing with contrastive explanations? Surprisingly he does not. His discussion is directed against the high probability theories of explanation. These theories accept that we can explain $f^*[c^*]$ in the case B, but not $f[c]$ in the case A. By referring to the arbitrariness of any probability threshold for explanation and by bringing in some symmetry considerations, he argues that the supporter of the high probability requirement should also accept explanations for the improbable outcomes. I take these arguments to be convincing. If we accept that we can explain f^* , we should also accept that we can explain f . And if we accept that we can explain $f^*[c^*]$, we should also accept that we can explain $f[c]$. However, we need some reasons to accept the antecedent of the counterfactual.

Salmon's justification seems to be that if we don't accept the antecedent, we cannot explain indeterministic events at all, and this would be too heavy a price to pay. Fortunately, we do not have to pay this price. The central intuition behind the contrastive approach to explanation is that the proper understanding of the explanation requires thinking carefully about the *explanandum*. Salmon is right that we can explain *something* about cases A and B. But the question is: what about them we can explain? And there are a lot of things to explain. For example, in case A, we can try explain:

- 1) why the probability of heads was .35 rather than some other value; or
- 2) why the result of the experiment was that the coin landed heads or tails, rather than something else, for example, that it melted, was vaporized, etc.

Question 1) directs us to the theoretically interesting case of explaining the propensity of the coin to land heads with some specific probability. Once we have an answer to this question, the

lack of an explanation for the outcome of a single experiment is much less important. Question 2) shows that there are many other contrasts to be explained, some of which are interesting. The only thing we cannot explain is $f[c]$, since it does not satisfy our requirements.

Is there anything else to explain? I don't think so. Contrary to Salmon, the contrastive approach is not a relic from the deterministic world that is only a hindrance to the proper understanding of explanation in the indeterministic contexts. Rather, it seems to be an indispensable tool for clarifying what we are capable of explaining in such contexts. Salmon agrees that we cannot completely explain f , but his theory does not allow him to explicate *what* about f can be explained. The contrastive theory can do this, which gives it a clear advantage.

And note how little we have to give up to be able to use this tool. We only have to give up the possibility of explaining singular outcomes of experiments (and actual frequencies that are the results of finite series of experiments), when both the fact and the foil have the possibility of occurring. And usually these are not the scientifically most interesting cases. For example, physicists working with quantum theory are not interested in explaining the results of single experiments or actual frequencies of outcomes in particular series of experiments. These actual frequencies are only evidence for the real objects of explanation, which are the probability distributions for the possible outcomes. (Watkins 1984: 228-229; Kitcher 1989: 450-452; Woodward 1989: 367-368.) Scientists are interested in the general phenomenon and its probability distributions, not in the idiosyncrasies of some specific set of data. This emphasizes the fact that we can have a deep theoretical understanding of the phenomenon without being able to predict or explain singular outcomes. The belief that the explanation of singular probabilistic outcomes is somehow scientifically central is an artefact of the philosophical discourse on explanation. The problem of the single case is philosophically challenging, but scientifically it is not a big deal.

In a later article, Salmon seems ready to concede this point. He accepts that he cannot convince his opponents about the importance of singular outcomes to theoretical basic science (Salmon 1998: 158). Without giving up his position, he tries to show that there are some cases of single-case explanations in the applied sciences that the above strategy cannot handle. He has basically two examples of such situations.

The first example concerns the outbreak of Legionnaires' disease in 1976. All the (original) victims of the disease had stayed in the same hotel, but not all hotel visitors contracted the disease. Actually, only a small percentage did. Later the bacillus responsible for the disease was found in the cooling towers for the air-conditioning system. Salmon suggests that since quantum fluctuations may lead to large uncertainties in the future trajectories of molecules in the air, we might not be able to provide a strictly deterministic explanation for why particular persons contracted the disease and some other people did not contract it in similarly infested rooms. Consequently, we cannot have a deductive explanation why precisely these people contracted the disease. However, he argues that for the purposes of assigning responsibility and

taking preventive steps in the future, we have an adequate explanation of the disease in this limited sample of people in 1976. (Salmon 1998: 158-159.)

The second example concerns eight soldiers (out of a group of 2235) who witnessed the detonation of an atomic bomb at close range during Operation Smoky in 1957, and who subsequently developed leukemia. According to Salmon, the high levels of radiation they were exposed to explain this incidence. Leukemia occurs with a non-zero frequency in the population at large, and it is possible, but not very likely, that these incidences of leukemia were due to a chance fluctuation rather than exposure to high levels of radiation. This makes the explanation probabilistic. Salmon argues that from a practical standpoint the fact to be explained is the high incidence of leukemia in this particular sample, not its incidence among all people who ever have been or will be exposed to that amount of radiation. Consequently, he argues, this is not a deductive statistical explanation. He also argues that this focus on this particular sample was of importance in the deciding whether the soldiers should receive extra compensation from the federal government. (Salmon 1998: 134, 159.)

Salmon's examples, and his arguments, are not convincing. (See Hällsten 1999, 2001 for a similar assessment.) In the case of Legionnaires' disease, we do not need to explain why these particular hotel guests contracted the disease in order to assign responsibility or to take preventive steps. For the latter, it is sufficient to know that there is a risk of Legionnaires' disease with certain kinds of air conditioning machines. It does not matter which individuals contracted the disease or exactly how many instances there were. The central issue is to eliminate the possibility of an outbreak of the disease. Similarly, when attributing responsibility, the significant issue is whether someone is responsible for putting the hotel visitors at risk, not the exact identity or the number of people contracting the disease. We don't have to explain why one particular guest contracted the disease rather than another. All we need to know is that it came from the same source. Consequently, we do not have to explain single cases, we only need to recognize that there is a number of incidences. Also, we can know the cause of their disease without being able to explain why they in particular contracted it.

Similar points apply to Salmon's leukemia example. The basis of the government's responsibility is in having put its the soldiers at a high risk, not causing particular soldiers to have leukemia (Kitcher 1989: 459). We do not have to explain the individual cases, we need note only their higher risk of getting this particular form of leukemia. And if a court decides that the risk was high enough, it can decide that the government is responsible, and that it must compensate all soldiers who were put at risk and who developed later leukemia. In this case we cannot know the cause of leukemia in individual cases (as we knew with Legionnaires' disease), but we know what increased their risk of having leukemia. This is enough for attributing legal and moral responsibility.

Let us review the position that I have adopted. The central idea is that a contrastive approach can be used to explicate what is explained and what is not. An approach that gives up the

Leibniz principle cannot have this advantage. The only drawback of this approach is that we cannot explain some singular outcomes of indeterministic processes. The Leibniz principle rules out cases where both f and c have a chance of occurrence after a occurs (within the intended causal field). However, we can explain $f[c]$ in situations where a and the presumed causal field do not necessitate f , but make it impossible for c to occur. This was the case in our earlier example of paresis. So we need only rule out the possibility of the occurrence of c , while still retaining the possibility that f will occur. We cannot predict the occurrence of f , but we can explain $f[c]$ after the fact.

We can use Mackie's famous example of three slot machines to illustrate how these ideas work in deterministic and indeterministic contexts. Mackie presents us with three different machines, K , L and M , which all profess to supply bars of chocolate and they all have a glass front that makes it possible to observe their internal mechanisms. Machine K is deterministic. That is what we expect of slot machines. It does not always produce a chocolate bar when a coin is inserted, but when it does not, we can always, in principle, find the fault in the machine or the outside interference that caused the failure to deliver. So, in normal circumstances, the insertion of a coin is both necessary and sufficient for the production of chocolate bar. Machine L is different. It is indeterministic. It will not, in normal circumstances, produce a chocolate bar when a coin is not inserted, but it may fail to produce a bar even when this is done. This failure is a matter of chance. We cannot explain the individual failures of L as we did in the case of K . So, in the normal circumstances, the insertion of a coin is only a necessary condition for the appearance of a chocolate bar. The third machine, M , is also indeterministic. In ordinary circumstances it will produce a bar of chocolate whenever a coin is inserted. However, it will also occasionally produce a chocolate bar even without insertion of a coin. In such cases the cause is not discoverable even in principle. So in a sense, M is the opposite of L , and for it the insertion of a coin is the sufficient condition of the production of a chocolate bar. (Mackie 1974: 40-41.)

We can explain why machine K produces a bar when a coin is inserted by describing the internal mechanisms that produce this result. Similarly we can explain why it does not produce a bar when a coin is not inserted. We can also explain the failures of the machine, because the glass front allows us to diagnose the failures. In the case of machine K we can always (in principle) explain the following *explananda* (among others):

- 1) the machine does not produce a bar when a coin is not inserted
[the machine produces a bar when a coin is not inserted].
- 2) the machine does not produce a bar when a coin is inserted
[the machine produces a bar when a coin is inserted].
- 3) the machine produces a bar when a coin is inserted
[the machine does not produce a bar when a coin is inserted].

4) the machine produces a bar when a coin is not inserted

[the machine does not produce a bar when a coin is not inserted].

With *L* and *M*, explaining all these *explananda* is not possible. In the case of *L*, we can explain 1) and 4), but we cannot explain 2) and 3). Note that we can explain 4) because the production of a bar would be a case of a deterministic mechanical failure observable through the glass front of the machine. Notice also that with machine *L* we can *ex post facto* know the cause of every production of a chocolate bar. It is either the insertion of a coin or a mechanical failure (or their combination). It is just that we cannot explain why a particular insertion of coin produces a bar and why some other insertion of a coin does not.

In the case of *M*, we can explain 2) and 3). Case 2) would again be a deterministic mechanical failure. However, in 3) there are two alternative mechanisms: the appearance of a bar might be a 'spontaneous' random event or it might be the result of the proper coin mechanism. So, in some cases we have an example of a causal preemption: the spontaneous mechanism preempts the production of the chocolate bar by the 'proper' mechanism. The explanation would have to mention this possible preemption, although the insertion of a coin ensures that a bar will be produced. But again, we have only two *explananda*, and we cannot explain 1) and 4).

Mackie's example does not mention any specific probability values for the indeterministic processes in *L* and *M*. Maybe there are no such things. But if there were, would it change the situation? My answer is that in the explanation of single occurrence they would not, no matter what the probability values are. Of course, the probabilities would help us in some other ways. They would help us in the theoretical description of the propensities inherent in the machines. It would also make it possible to predict outcomes more accurately. But they would not affect the explanation.

Note that the probabilities discussed above reflect truly indeterministic processes. However, of the most probabilities used in the sciences and in everyday life are not like this. These probabilities reflect our ignorance, and my points do not apply to them. In these cases we have only *partial* explanations. One advantage of my position is that it does not make these explanations look better than they are. The existence of indeterministic phenomena should not lead us to be more satisfied with our usual statistical explanations as they deserve. Some accounts of probabilistic explanation do have this vice, as they accept a partial explanation as a fully adequate explanation.

I conclude that indeterministic causal processes do not produce any fatal problems for my account of causal explanation. I hope that I have also satisfied Salmon, who urges deductivists to consider cases like Mackie's chocolate machines (Salmon 1989: 176-177).

2. Dispensing with probabilistic versions of contrastive theory

One approach that I have not considered so far is the suggestion that we should understand contrastive explanation in probabilistic terms. There are slightly different versions of this idea in the published literature. Although they differ in the implementation of this basic idea, they all propose that the explanatory cause makes f more probable than c .

The first version is van Fraassen's suggestion that A should favor F against C . He interprets this favoring to mean that the A raises the probability of F in comparison to C (van Fraassen 1980: 146-147).² Twenty years later Philip Percival suggested that we should require that A strongly favors F against C . Although his proposal was not fully developed, it seems that A strongly favors F against C when it makes the probability of F near one and the probability of c near zero (Percival 2000: 63). Percival's position is consequently more demanding than van Fraassen's. Elliot Sober, on the other hand, has taken a more liberal stance. His suggestion is that a common cause *discriminates* between F and C . For him, this means that the common cause makes F more probable than C . (Sober 1994: 178.)

There is also an interesting suggestion by Christopher Hitchcock (1999) that does not require that the explanatory cause makes F probable or that the explanatory factor raises the probability of F . In this account it is only required that the explanatory factor make a difference in the probability of F against the probability of C . I will return to Hitchcock's model after considering the ideas presented in earlier theories.

I think that there are some weighty general arguments against the whole idea of a probabilistic approach to contrastive explanation. I will not go to the details of the suggestions made by van Fraassen, Sober and Percival. Instead I will go directly to the arguments against the underlying idea.

The first argument is actually a question. These probabilistic theories of contrastive explanation start from the idea that an explanation makes the *explanandum* expected. This is the old Hempelian intuition about explanation. Despite this influential source, I think we should ask *what explanation has to do with expectedness*. Once we have given up the idea about the symmetry between explanation and prediction (or rational expectability) there is no positive reason to entertain this connection. If explanation is really about answering why- and how-questions, the connection between explanation and high probabilities seems arbitrary. Explanation deals principally with mechanisms and processes, not with the probabilities of the various outcomes. It might be that some philosophers have a contrary intuition, but we should ask whether their intuitions are simply products of too much exposure to the D-N model? Of course, the causal origin of the idea does not invalidate it, but if we have no other justification for it, we should consider whether we should still entertain it. What we need is some positive motive to connect statistical relevance with explanation. The authors under consideration here have not offered any.

² This can be understood for example as: $\Pr(F|A) > \Pr(F) \ \& \ \Pr(C|A) < \Pr(C)$ & $\Pr(F|A) > \Pr(C|A)$. However, other reconstructions are also possible. In the following, letters refer to the types, small letters to the tokens.

The second argument is by Wesley Salmon. Let us call it *the symmetry argument*. Recall our experimental set-up from example 1. The theories we are now considering suggest that we can explain the coin's landing heads in case B, where the switch was turned up, but we cannot explain its landing heads in case A, where the switch was turned down, whereas if the coin had turned tails, we could have explained it in case A, but not in case B. Now Salmon's symmetry argument asks why we could not explain the heads in case A, if we can explain the tails if it were the result of the experiment? After all, the causal facts are the same. The components and mechanisms of the causal process are exactly the same in both cases, only the outcome is different. If the explanation is about understanding causal processes and mechanisms, as suggested in the previous argument, there should not be any difference in the explanations of f [c] and c [f]. We understand the causal process equally well in both processes, so where is the explanatory difference? (Salmon 1998: 152-153.)

The third argument is also by Wesley Salmon. Let us call it *the arbitrariness argument*. Salmon notes that a supporter of the high probability requirement for explanation has trouble with justifying any specific probability threshold for explanatoriness. Suppose that we require that A must make F more probable than C and that we are comparing the following two cases. In the first case the probability of F is .501 and in the second case its probability is .499. According to the view under consideration we should judge that A is explanatory in the first case, but not in the second. This judgment feels arbitrary, since the difference between the probabilities is so small. We can develop a similar argument for any possible probability threshold (Salmon 1998: 151-152). The differences among the suggestions of van Fraassen, Sober, and Percival reflect the problem: they all characterize probability threshold slightly differently. Why? No justification is given. Maybe there is no such thing, and the whole idea of the high probability requirement for explanation is misconceived.

We have still the fourth and, to my mind, the strongest argument. The idea for this argument comes from Woodward (1990), who argues that there are some difficult problems for current theories of causation in the context of applying probabilities to singular causal statements. The basic idea of this argument is simply that probabilities do not reflect facts about singular causal processes, as they should. Consider the following example:

Example 2: A recovery with multiple causes

Suppose there is a disease that has a spontaneous recovery rate of .05 within two weeks. There are two alternative treatments, which work by completely different causal mechanisms from each other and from spontaneous recovery. There are also no interaction effects between the treatments. Treatment A raises the probability of recovery within two weeks by .47 and treatment B raises the same probability by .28. Both treatments are used to treat a patient who recovers within two weeks. Our explanation-seeking question is the following: Why did the patient recover within two weeks rather than later?

Let us now consider the causal facts of the case. We have at least seven alternative causal histories for the recovery of our patient:

- 1) the recovery was caused by a (spontaneous recovery)
- 2) the recovery was caused by a^* (treatment A)
- 3) the recovery was caused by a^{**} (treatment B)
- 4) the recovery was caused by the combination a & a^*
- 5) the recovery was caused by the combination a & a^{**}
- 6) the recovery was caused by the combination a^* & a^{**}
- 7) the recovery was caused by the combination a & a^* & a^{**} .

There are alternative ways of interpreting the combination of causes in cases 4) - 7). The combination of causes can be understood either as causal overdetermination, causal pre-emption, or joint causation. Consequently, we can have up to 20 possible alternative causal histories.

However, from the point of view of my argument, the exact number of alternatives does not matter. What is important is that *the probabilistic theories of explanation rule out most of these causal possibilities on purely conceptual grounds*. Intuitively we think that the explanation should reflect the causal facts of the case. So, if the recovery was caused by treatment B, then the explanation of the recovery is the successful use of treatment B. Similarly, if the recovery was spontaneous, we should not claim that the treatments were effective in this particular recovery. Probabilistic theories do not allow us this kind of sensitivity to the facts. We have to accept 2) (or alternatively 7)) as the right explanation since it makes the recovery most probable. This judgment is not affected by what happens in this particular case. It might be that different probabilistic theories will pick different causal factors (or their combinations) as explanatory, but none of them allows all seven (or twenty) explanatory scenarios. The theories accept that the events might have unfolded according to any of the above causal scenarios, but only one of them can be presented as an explanation, independently of its truth in this particular case. This is clearly unacceptable: why should statistics about other cases determine what happens in one particular case?

I think that the correct judgment is that with the information provided we cannot yet explain why the patient recovered within two weeks. All causal scenarios should be kept open. What is needed is a closer examination of the patient's recovery process in order to find out how recovery occurred. It might be possible to rule out some of the alternative causes. Of course, it is also possible that we would never find out how recovery occurred. In such a case, we should notice the causal possibilities (and their probabilities) and remain uncommitted to any claims about the true explanation. There is a fact of the matter, but we just do not know it. Theories of explanation should respect this possibility rather than prejudging the issue in the manner of the

probabilistic theories under consideration.

A similar lesson can be drawn from other examples, where the actual cause either diminishes the probability of the *explanandum* or it does not make any difference in its probability. (For a wide array of such examples, see Eells 1991.) The crucial point is that the probabilistic accounts do not treat causal facts about the causal history correctly, and they should fail because of this shortcoming. The singular causal explanation should reflect the facts of the particular case, not what usually happens. Similarly, if the causal process producing the event to be explained includes complications like overdetermination, joint causes or causes that diminish the probability of the effect, the explanation should spell out these complexities instead of covering them up with probability statements.

My own suggestion is that we treat probabilistic cases as partial explanations. They provide us part of the explanation without fully explaining the intended contrast. This would allow us to treat probabilistic information as contributing to explanations without compromising our standards for them. This is important because most of our probabilistic explanations reflect more of our ignorance than the causal indeterminacy of the world. In lowering standards by adopting a probabilistic account of explanation, one only fools oneself.

Hitchcock's explanatory relevance model

Christopher Hitchcock (1999) has presented a theory that avoids the first three critiques presented above. The fourth argument applies to his theory, and I consider it to be enough. However, since Hitchcock's paper adds some new elements to the literature on contrastive explanation, I think it is appropriate to comment on some of his ideas and to clarify my own position against them.

Hitchcock argues that we can explain probabilistically why *F* rather than *E* happened. His basic idea is as follows (*B* refers to the background conditions):

... we should view *A* as explanatorily relevant to the contrastive question 'why *E* rather than *F* if *A* continues to be relevant to *E* when the (exclusive) disjunction $E \vee F$ is held fixed. [...] this means that *A* is explanatorily relevant to *E* rather than *F* when $P(E | (A \& B) \& (E \vee F)) > P(E | B \& (E \vee F))$ (Hitchcock 1999: 587).

Now, the most important change that Hitchcock makes to the settings of the earlier discussion is that he speaks of *explanatory relevance* instead of *explanation*. This change is very significant. According to the Leibniz principle, *a* cannot *explain* (in similar circumstances) both *f* and *c* if they are incompatible with each other. However, the principle says nothing about *a*'s being explanatorily relevant.

Let us have a look at Hitchcock's central example to see how his idea works. Suppose a vertically polarized photon approaches a polarizer that is tilted at an angle from the vertical. Now this is, according to our current knowledge, a chance set-up. There is a certain chance the photon will

pass through the polarizer and a certain chance that it will be absorbed by it. (Hitchcock 1999: 585-586.) Let us suppose that that the photon was transmitted. Now the Leibniz principle would say that we cannot explain why the photon was transmitted rather than absorbed. Hitchcock wishes to disagree. His motivation becomes understandable when we add more information to the example. The fact to be added is that the probability of transmission varies with the relative orientation of the photon and the polarizer. Depending on the angle of contact the probability varies between zero and one. Hitchcock's point is that the relative orientation of photon and polarizer is explanatorily relevant because it affects the probability of the transmission. (Hitchcock 1999: 596.)

Now, I agree with Hitchcock that the relative orientation is explanatorily relevant. However, the big question is: *for which explanandum it is explanatorily relevant?* When we ask why the transmission had a certain probability rather than some other probability, the reference to the relative orientations of the photon and the polarizer is certainly appropriate. After all, the angle makes the difference between probabilities of transmission. In other words, it *explains* the probability in question (in contrast to some alternatives). The possible controversy concerns the explanation why the photon was transmitted rather than absorbed. Again, nobody would deny that the relative orientation of the photon and the polarizer was a *causally* relevant factor in the process of the transmission of the photon. There is no disagreement about facts of the example. The issue concerns the explanation of a singular event.

Now, let us suppose that someone asks why a particular photon was transmitted rather than absorbed? Would there be some differences between the explanatory answers given by Hitchcock and by more conservative theorists? The answer Hitchcock would give to this question would refer to the probabilistic nature of the process and to the role of the relative orientation of the photon and the polarizer in this process. After all, there is some causally relevant information to be provided. How would Hitchcock's opponent answer this question? I do not think that she would provide 'nothing' as her first reply to this question, as Hitchcock seems to imply. She would probably give exactly the same answer to the question, but she would add that these facts are not enough to explain the contrast. So with this explanation-seeking question we cannot establish the real difference between the positions. Now let us move to the follow-up question: *given that the transmission had a certain probability to occur, why was the photon transmitted?* I think that the both parties would give exactly the same answer here: 'we can't explain that, because it was due to chance'. These considerations seem to suggest that there is no real disagreement between Hitchcock and me as to how to handle this example.

Since Hitchcock cannot show that any real advantage derives from his proposal, I think we should not give up Leibniz's rather intuitive principle. Rejecting it might lead us to situations where explanation comes too easy and is intuitively non-satisfactory. Hitchcock might reply to this by claiming that his opponents cannot really do justice to explanations that do not fit their

ideal standards. By this he would point to the fact that the intuitive answer to the above example is not really permitted by his opponents' standards of explanation. I think this argument is a non-starter. Hitchcock's points can be easily accommodated to the more traditional account without giving up the Leibniz principle.

First, we should distinguish between an explanation of an event and an explanation of that event's probability. Hitchcock might be confusing two different *explananda*. When explaining the probability of the individual event, Hitchcock's considerations are relevant. But they are not relevant when we are explaining a singular occurrence. And this is the only case where his considerations would contradict the received view. When we are explaining a singular occurrence, which is, as I have already noted, of peripheral scientific interest, the thing either happens or it does not. This is all there is to it. I cannot see how the probability of this occurrence could have any causal influence here. Suggesting otherwise would imply that chance is one causal factor among others, and this would be a category mistake (Humphreys 1989).

Second, despite the fact that I set high standards for explanatory adequacy, it does not follow that I do not see the value of partial explanation. Sometimes a more detailed description of the process leading to the event to be explained is valuable without being able to explain the contrast we have in mind. Sometimes this description can also work as a correction to the question with an inadequate assumption. A partial explanation is better than no explanation at all, so I think Hitchcock's worries can be handled.

References

- Hitchcock, Christopher C. 1999: 'Contrastive Explanation and the Demons of Determinism', *The British Journal for the Philosophy of Science* 50: 585-612.
- Hållsten, Henrik 1999: 'Deductive Chauvinism', *Synthese* 120: 49-59.
- Hållsten, Henrik 2001: *Explanation and Deduction. A Defence of Deductive Chauvinism*. Acta Universitatis Stockholmiensis. Stockholm Studies on Philosophy 21.
- Kitcher, Philip 1989: 'Explanatory Unification and the Causal Structure of the World', in Kitcher & Salmon (eds.): *Scientific Explanation. Minnesota Studies in the Philosophy of Science vol XIII*. University of Minnesota Press. Minneapolis: 410-505.
- Lewis, David 1986: *Philosophical Papers vol II*. Oxford University Press. Oxford.
- Mackie, J. L. 1974: *The Cement of the Universe*. Clarendon Press. Oxford.
- Percival, Philip 2000: 'Lewis's Dilemma of Explanation under Indeterminism Exposed and Resolved', *Mind* 109: 39-66.
- Ruben, David-Hillel 1990: *Explaining Explanation*. Routledge. London.
- Salmon, Wesley 1998: *Causality and Explanation*. Oxford University Press. Oxford.
- Sober, Elliot 1994: *From a Biological Point of View*. Cambridge University Press. Cambridge.
- van Fraassen, Bas 1980: *Scientific Image*. Oxford University Press. Oxford.
- Watkins, John 1984: *Science and Scepticism*. Princeton University Press. Princeton.
- Woodward, James 1989: 'The Causal Mechanical Model of Explanation', in Kitcher & Salmon

(eds.): *Scientific Explanation. Minnesota Studies in the Philosophy of Science vol XIII*. University of Minnesota Press. Minneapolis: 357-383.

Woodward, James 1990: 'Supervenience and Singular Causal Claims', in D. Knowles (ed.) *Explanation and its limits*. Cambridge University Press. Cambridge.

Ylikoski, Petri 2001: *Understanding Interests and Causal Explanation*. Ph.D.-thesis, available at <http://ethesis.helsinki.fi/julkaisut/val/kayta/vk/ylikoski/>