The American Philosophical Association
University of Delaware
Newark, DE 19716

Why Ask, "Why?"? An Inquiry concerning Scientific Explanation
Author(s): Wesley C. Salmon
Source: *Proceedings and Addresses of the American Philosophical Association*, Vol. 51, No. 6 (Aug., 1978), pp. 683-705
Published by: American Philosophical Association
Stable URL: http://www.jstor.org/stable/3129654
Accessed: 24/08/2013 12:32

WHY ASK, "WHY?"?

AN INQUIRY CONCERNING SCIENTIFIC EXPLANATION

Wesley C. Salmon[1]
University of Arizona

Concerning the first order question "Why?" I have raised the second order question "Why ask, 'Why?'?" to which you might naturally respond with the third order question "Why ask, 'Why ask, "Why?'?" But this way lies madness, to say nothing of an infinite regress. While an infinite sequence of nested intervals may converge upon a point, the series of nested questions just initiated has no point to it, and so we had better cut it off without delay. The answer to the very natural third order question is this: the question "Why ask, 'Why?'?" expresses a deep philosophical perplexity which I believe to be both significant in its own right and highly relevant to certain current philosophical discussions. I want to share it with you.

The problems I shall be discussing pertain mainly to scientific explanation, but before turning to them, I should remark that I am fully aware that many--perhaps--most-why-questions are requests for some sort of justification (Why did one employee receive a larger raise than another? – Because she had been paid less than a male colleague for doing the same kind of job.) or consolation (Why, asked Job, was I singled out for such extraordinary misfortune and suffering?). Since I have neither the time nor the talent to deal with questions of this sort, I shall not pursue them further, except to remark that the seeds of endless philosophical confusion can be sown by failing carefully to distinguish them from requests for scientific explanation.

---

*Presidential Address delivered before the Fifty-second Annual Pacific Meeting of the American Philosophical Association in San Francisco, California, March 24, 1978.

Let me put the question I do want to discuss to you this way. Suppose you had achieved the epistemic status of Laplace's demon – the hypothetical super-intelligence who knows all of nature's regularities, and the precise state of the universe in full detail at some particular moment (say now, according to some suitable simultaneity slice of the universe). Possessing the requisite logical and mathematical skill, you would be able to predict any future occurrence, and you would be able to retrodict any past event. Given this sort of apparent omniscience, would your scientific knowledge be complete, or would it still leave something to be desired? Laplace asked no more of his demon; should we place further demands upon ourselves? And if so, what should be the nature of the additional demands?

If we look at most contemporary philosophy of science texts, we find an immediate affirmative answer to this question. Science, the majority say, has at least two principal aims--prediction (construed broadly enough to include inference from the observed to the unobserved, regardless of temporal relations) and explanation. The first of these provides knowledge of what happens; the second is supposed to furnish knowledge of why things happen as they do. This is not a new idea. In the Posterior Analytics, Aristotle distinguishes syllogisms which provide scientific understanding from those which do not.[2] In the Port Royal Logic, Arnauld distinguishes demonstrations which merely convince the mind from those which also enlighten the mind.[3]

This view has not been universally adopted. It was not long ago that we often heard statements to the effect that the business of science is to predict, not to explain. Scientific knowledge is descriptive–it tells us what and how. If we seek explanations–if we want to know why–we must go outside of science, perhaps to metaphysics or theology. In his Preface to the Third Edition (1911) of The Grammar of Science, Karl Pearson wrote, "Nobody believes now that science explains anything; we all look upon it as a shorthand description, as an economy of thought."[4] This doctrine is not very popular nowadays. It is now fashionable to say that science aims not merely at describing the world--it also provides understanding, comprehension, and enlightenment. Science presumably accomplishes such high-sounding goals by supplying scientific explanations.

The current attitude leaves us with a deep and perplexing question, namely, if explanation does involve something over and above mere description, just what sort of thing is it? The use of such honorific near-synonyms as "understanding," "comprehension," and "enlightenment" makes it sound important and desirable, but does not help at all in the philosophical analysis of explanation--scientific or other. What, over and above its complete descriptive knowledge of the world, would Laplace's demon require in order to achieve understanding? I hope you can see that this is a real problem, especially for those who hold what I shall call "the inferential view" of scientific explanation, for Laplace's demon can infer every fact about the universe, past, present, and future. If the problem does not seem acute, I would quote a remark made by Russell about Zeno's paradox of the flying arrow--"The more the difficulty is meditated, the more real it becomes."[5]

It is not my intention this evening to discuss the details of the various formal models of scientific explanation which have been advanced in the last three decades.[6] Instead, I want to consider the general conceptions which lie beneath the most influential theories of scientific explanation. Two powerful intuitions seem to have guided much of the discussion. Although they have given rise to disparate basic conceptions and considerable controversy, both are, in my opinion, quite sound. Moreover, it seems to me, both can be incorporated into a single overall theory of scientific explanation.

(1) The first of these intuitions is the notion that the explanation of a phenomenon essentially involves locating and identifying its cause or causes. This intuition seems to arise rather directly from common sense, and from various contexts in which scientific knowledge is applied to concrete situations. It is strongly supported by a number of paradigms, the most convincing of which are explanations of particular occurrences. To explain a given airplane crash, for example, we seek “the cause”--a mechanical failure, perhaps, or pilot error. To explain a person’s death again we seek the cause--strangulation or drowning, for instance. I shall call the general view of scientific explanation which comes more or less directly from this fundamental intuition the causal conception; Michael Scriven has been one of its chief advocates.[7]

(2) The second of these basic intuitions is the notion that all scientific explanation involves subsumption under laws. This intuition seems to arise from consideration of developments in theoretical science. It has led to the general “covering law” conception of explanation, as well as to several formal “models” of explanation. According to this view, a fact is subsumed under one or more general laws if the assertion of its occurrence follows, either deductively or inductively, from statements of the laws (in conjunction, in some cases, with other premises). Since this view takes explanations to be arguments, I shall call it the inferential conception; Carl G. Hempel has been one of its ablest champions.[8]

Although the proponents of this inferential conception have often chosen to illustrate it with explanations of particular occurrences--e.g., why did the bunsen flame turn yellow on this particular occasion?–the paradigms which give it strongest support are explanations of general regularities. When we look to the history of science for the most outstanding cases of scientific explanations, such examples as Newton’s explanation of Kepler’s laws of planetary motion or Maxwell’s electromagnetic explanation of optical phenomena come immediately to mind.

It is easy to guess how Laplace might have reacted to my question about his demon, and to the two basic intuitions I have just mentioned. The super-intelligence would have everything needed to provide scientific explanations. When, to mention one of Laplace’s favorite examples, a seemingly haphazard phenomenon, such as the appearance of a comet, occurs, it can be explained by showing that it actually conforms to natural laws.[9] On Laplace’s assumption of determinism, the demon possesses explanations of all happenings in the entire history of the world--past, present, and future. Explanation, for

Laplace, seemed to consist in showing how events conform to the laws of nature, and these very laws provide the causal connections among the various states of the world. The Laplacian version of explanation thus seems to conform both to the causal conception and to the inferential conception.

Why, you might well ask, is not the Laplacian view of scientific explanation basically sound? Why do twentieth century philosophers find it necessary to engage in lengthy disputes over this matter? There are, I think, three fundamental reasons: (1) the causal conception faces the difficulty that no adequate treatment of causation has yet been offered; (2) the inferential conception suffers from the fact that it seriously misconstrues the nature of subsumption under laws; and (3) both conceptions have overlooked a central explanatory principle.

The inferential view, as elaborated in detail by Hempel and others, has been the dominant theory of scientific explanation in recent years--indeed, it has become virtually "the received view." From that standpoint, anyone who had attained the epistemic status of Laplace's demon could use the known laws and initial conditions to predict a future event, and when the event comes to pass, the argument which enabled us to predict it would ipso facto constitute an explanation of it. If, as Laplace believed, determinism is true, then every <u>future</u> event would thus be amenable to deductive-nomological explanation.

When, however, we consider the explanation of past events--events which occurred earlier than our initial conditions--we find a strange disparity. Although, by applying known laws, we can reliably <u>retrodict</u> any past occurrence on the basis of facts subsequent to the event, our intuitions rebel at the idea that we can <u>explain</u> events in terms of subsequent conditions. Thus, although our inferences to future events qualify as explanations according to the inferential conception, our inferences to the past do not. Laplace's demon can, of course, construct explanations of past events by inferring the existence of still earlier conditions and, with the aid of the known laws, deducing the occurrence of the events to be explained from these conditions which held in the more remote past. But if, as the inferential conception maintains, explanations are essentially inferences, such an approach to explanation of past events seems strangely roundabout. Explanations demand an asymmetry not present in inferences.

When we drop the fiction of Laplace's demon, and relinquish the assumption of determinism, the asymmetry becomes even more striking. The demon can predict the future and retrodict the past with complete precision and reliability. We cannot. When we consider the comparative difficulty of prediction vs. retrodiction, it turns out that retrodiction enjoys a tremendous advantage. We have records of the past–tree rings, diaries, fossils–but none of the future. As a result, we can have extensive and detailed knowledge of the past which has no counterpart in knowledge about the future. From a newspaper account of an accident, we can retrodict all sorts of details which could not have been predicted an hour before the collision. But the newspaper story–even though it may *report* the explanation of the accident--surely does not *constitute* the explanation. We see that *inference* has a preferred temporal direction, and that *explanation* also has a preferred temporal direction. The fact that these two are opposite to each other is one thing which makes me seriously doubt that explanations are essentially arguments.[10] As we shall see, however, denying that explanations are arguments does not mean that we must give up the *covering law* conception. Subsumption under laws can take a different form.

Although the Laplacian conception bears strong similarities to the received view, there is a fundamental difference which must be noted. Laplace evidently believed that the explanations provided by his demon would be *causal explanations,* and the laws invoked would be *causal laws.* Hempel's deductive-nomological explanations are often casually called "causal explanations," but this is not accurate.[11] Hempel explicitly notes that some laws, such as the ideal gas law,

$$PV = nRT,$$

are non-causal. This law states a mathematical functional relationship among several quantities--pressure P, volume V, temperature T, number of moles of gas n, universal gas constant R--but gives no hint as to how a change in one of the values would lead causally to changes in others. As far as I know, Laplace did not make any distinction between causal and non-causal laws; Hempel has recognized the difference, but he allows non-causal as well as causal laws to function as covering laws in scientific explanations.

This attitude toward non-causal laws is surely too tolerant. If someone inflates an air-mattress of a given size to a certain pressure under conditions which determine the temperature, we can deduce the value of n--the amount of air blown into it. The *subsequent* values of pressure, temperature, and volume are thus taken to explain the quantity of air *previously* introduced. Failure to require covering laws to be causal laws leads to a violation of the temporal requirement on explanations. This is not surprising. The asymmetry of explanation is inherited from the asymmetry of causation--namely, that causes precede their effects. At this point, it seems to me, we experience vividly the force of the intuitions underlying the causal conception of scientific explanation.

There is another reason for maintaining that non-causal laws cannot bear the burden of covering laws in scientific explanations. Non-causal regularities, instead of having explanatory force which enables them to provide understanding of events in the world, cry out to be explained. Mariners, long before Newton, were fully aware of the correlation between the behavior of the tides and the position and phase of the moon. But inasmuch as they were totally ignorant of the causal relations involved, they rightly believed that they did not understand why the tides ebb and flow. When Newton provided the gravitational links, understanding was achieved. Similarly, I should say, the ideal gas law had little or no explanatory power until its causal underpinnings were furnished by the molecular-kinetic theory of gases. Keeping this consideration in mind, we realize that we must give at least as much attention to the explanations of regularities as we do to explanations of particular facts. I will argue, moreover, that these regularities demand causal explanation. Again, we must give the causal conception its due.

Having considered a number of preliminaries, I should now like to turn my attention to an attempt to outline a general theory of causal explanation. I shall not be trying to articulate a formal model; I shall be focusing upon general conceptions and fundamental principles rather than technical details. I am not suggesting, of course, that the technical details are dispensable--merely that this is not the time or place to try to go into them. Let me say at the outset that I shall be relying very heavily upon works by Russell (especially, The Analysis of Matter and Human Knowledge, Its Scope and Limits) and Reichenbach (especially, The Direction of Time). Although, to the best of my knowledge, neither of these authors ever published an article, or a book, or a chapter of a book devoted explicitly to scientific explanation, nevertheless, it seems to me that a rather appealing theory of causal explanation can be constructed by putting together the insights expressed in the aforementioned works.

Developments in twentieth-century science should prepare us for the eventuality that some of our scientific explanations will have to be statistical--not merely because our knowledge is incomplete (as Laplace would have maintained), but rather, because nature itself is inherently statistical. Some of the laws used in explaining particular events will be statistical, and some of the regularities we wish to explain will also be statistical. I have been urging that causal considerations play a crucial role in explanation; indeed, I have just said that regularities--and this certainly includes statistical regularities--require causal explanation. I do not believe there is any conflict here. It seems to me that, by employing a statistical conception of causation along the lines developed by Patrick Suppes and Hans Reichenbach,[12] it is possible to fit together harmoniously the causal and statistical factors in explanatory contexts. Let me attempt to illustrate this point by discussing a concrete example.

A good deal of attention has recently been given in the press to cases of leukemia in military personnel who witnessed an atomic bomb test (code name "Smokey") at close range in 1957.[13] Statistical studies of the survivors of the bombings of Hiroshima and Nagasaki have established the fact that exposure to high levels of radiation, such as occur in an atomic blast, is statistically relevant to the occurrence of leukemia--indeed, that the probability of leukemia is closely correlated with the distance from the explosion.[14] A clear pattern of statistical relevance relations is exhibited here. If a particular person contracts leukemia, this fact may be explained by citing the fact that he was, say, 2 kilometers from the hypocenter at the time of the explosion. This relationship is further explained by the fact that individuals located at specific distances from atomic blasts of specified magnitude receive certain high doses of radiation.

This tragic example has several features to which I should like to call special attention:

(1) The location of the individual at the time of the blast is statistically relevant to the occurrence of leukemia; the probability of leukemia for a person located 2 kilometers from the hypocenter of an atomic blast is radically different from the probability of the disease in the population at large. Notice that the probability of such an individual contracting leukemia is not high; it is much smaller than one-half--indeed, in the case of Smokey it is much less than 1/100. But it is markedly higher than for a random member of the entire human population. It is the *statistical relevance* of exposure to an atomic blast, not a high probability, which has explanatory force.[15] Such examples defy explanation according to an inferential view which requires high inductive probability for statistical explanation.[16] The case of leukemia is subsumed under a statistical regularity, but it does not “follow inductively” from the explanatory facts.

(2) There is a *causal process* which connects the occurrence of the bomb blast with the physiological harm done to people at some distance from the explosion. High energy radiation, released in the nuclear reactions, traverses the space between the blast and the individual. Although some of the details may not yet be known, it is a well-established fact that such radiation does interact with cells in a way which makes them susceptible to leukemia at some later time.

(3) At each end of the causal process-i.e., the transmission of radiation from the bomb to the person--there is a *causal interaction.* The radiation is emitted as a result of a nuclear interaction when the bomb explodes, and it is absorbed by cells in the body of the victim. Each of these interactions may well be irreducibly statistical and indeterministic, but that is no reason to deny that they are causal.

(4) The causal processes begin at a central place, and they travel outward at a finite velocity. A rather complex set of statistical relevance relations is explained by the propagation of a process, or set of processes, from a common central event.

In undertaking a general characterization of causal explanation, we must begin by carefully distinguishing between causal processes and causal interactions. The transmission of light from one place to another, and the motion of a material particle, are obvious examples of causal processes. The collision of two billiard balls, and the emission or absorption of a photon, are standard examples of causal interactions. Interactions are the sorts of things we are inclined to identify as events. Relative to a particular context, an event is comparatively small in its spatial and temporal dimensions; processes typically have much larger durations, and they may be more extended in space as well.

A light ray, traveling to earth from a distant star, is a process which covers a large distance and lasts for a long time. What I am calling a "causal process" is similar to what Russell called a "causal line."[17]

When we attempt to identify causal processes, it is of crucial importance to distinguish them from such pseudo-processes as a shadow moving across the landscape. This can best be done, I believe, by invoking Reichenbach's *mark criterion.*[18] Causal processes are capable of propagating marks or modifications imposed upon them; pseudo-processes are not. An automobile traveling along a road is an example of a causal process. If a fender is scraped as a result of a collision with a stone wall, the mark of that collision will be carried on by the car long after the interaction with the wall occurred. The shadow of a car moving along the shoulder is a pseudo-process. If it is deformed as it encounters a stone wall, it will immediately resume its former shape as soon as it passes by the wall. It will not transmit a mark or modification. For this reason, we say that a causal process can transmit information or causal influence; a pseudo-process cannot.[19]

When I say that a causal process has the capability of transmitting a causal influence, it might be supposed that I am introducing precisely the sort of mysterious power Hume warned us against. It seems to me that this danger can be circumvented by employing an adaptation of the "at-at" theory of motion, which Russell used so effectively in dealing with Zeno's paradox of the flying arrow.[20] The flying arrow--which is, by the way, a causal process--gets from one place to another by being *at* the appropriate intermediate points of space *at* the appropriate instants of time. Nothing more is involved in getting *from* one point *to* another. A mark, analogously, can be said to be propagated from the point of interaction at which it is imposed to later stages in the process if it appears *at* the appropriate intermediate stages in the process *at* the appropriate times without additional interactions which regenerate the mark. The precise formulation of this condition is a bit tricky, but I believe the basic idea is simple, and that the details can be worked out.[21]

If this analysis of causal processes is satisfactory, we have an answer to the question, raised by Hume, concerning the connection between cause and effect. If we think of a cause as one event, and of an effect as a distinct event, then the connection between them is simply a spatio-temporally continuous causal process. This sort of answer did not occur to Hume because he did not distinguish between causal processes and causal interactions. When he tried to analyze the connections between distinct events, he treated them as if they were chains of events with discrete links, rather than processes analogous to continuous filaments. I am inclined to attribute considerable philosophical significance to the fact that each link in a chain has adjacent links, while the points in a continuum do not have next-door neighbors. This consideration played an important role in Russell's discussion of Zeno's paradoxes.[22]

After distinguishing between causal interactions and causal processes, and after introducing a criterion by means of which to discriminate the pseudo-processes from the genuine causal processes, we must consider certain configurations of processes which have special explanatory import. Russell noted that we often find similar structures grouped symmetrically about a center--for example, concentric waves moving across an otherwise smooth surface of a pond, or sound waves moving out from a central region, or perceptions of many people viewing a stage from different seats in a theatre. In such cases, Russell postulates the existence of a central event--a pebble dropped into the pond, a starter's gun going off at a race-track, or a play being performed upon the stage--from which the complex array emanates.[23] It is noteworthy that Russell never suggests that the central event is to be explained on the basis of convergence of influences from remote regions upon that locale.

Reichenbach articulated a closely-related idea in his <u>principle of the common cause</u>. If two or more events of certain types occur at different places, but occur at the same time more frequently than is to be expected if they occurred independently, then this apparent coincidence is to be explained in terms of a common causal antecedent.[24] If, for example, all of the electric lights in a particular area go out simultaneously, we do not believe that they just happened by chance to burn out at the same time. We attribute the coincidence to a common cause such as a blown fuse, a downed transmission line, or trouble at the generating station. If all of the students in a dormitory fall ill on the same night, it is attributed to spoiled food in the meal which all of them ate. Russell's similar structures arranged symmetrically about a center obviously qualify as the sorts of coincidences which require common causes for their explanations.[25]

In order to formulate his common cause principle more precisely , Reichenbach defined what he called a <u>conjunctive fork</u>. Suppose we have events of two types, <u>A</u> and <u>B</u>, which happen in conjunction more often than they would if they were statistically independent of one another. For example, let <u>A</u> and <u>B</u> stand for colorblindness in two brothers. There is a certain probability that a male, selected from the population at random, will have that affliction, but since it is hereditary, occurrences in male siblings are not independent. The probability that both will have it is greater than the product of the two respective probabilities. In cases of such statistical dependencies, we invoke a common cause <u>C</u> which accounts for them; in this case, it is a genetic factor carried by the mother. In order to satisfy the conditions for a conjunctive fork, events of the types <u>A</u> and <u>B</u> must occur independently in the absence of the common cause <u>C</u>--that is, for two unrelated males, the probability of both being colorblind is equal to the product of the two separate probabilities. Furthermore, the probabilities of <u>A</u> and <u>B</u> must each be increased above their overall values if <u>C</u> is present. Clearly the probability of colorblindness is greater in sons of mothers carrying the genetic factor than it is among all male children regardless of the genetic make-up of their mothers. Finally, Reichenbach stipulates, the dependency between <u>A</u> and <u>B</u> is absorbed into the occurrence of the common cause <u>C</u>, in the sense that the probability of <u>A</u> and <u>B</u> given <u>C</u> equals the product of the probability of <u>A</u> given <u>C</u> and the probability of B given C. This is true in the colorblindness case. Excluding pairs of identical twins, the question of whether a male child inherits colorblindness from the mother who carries the genetic trait depends only upon the genetic relationship between that child and his mother, not upon whether other sons happened to inherit the trait.[26] Note that screening-off occurs here.[27] While the colorblindness of a brother is statistically relevant to colorblindness in a boy, it becomes irrelevant if the genetic factor is known to be present in the mother.

Reichenbach obviously was not the first philosopher to notice that we explain coincidences in terms of common causal antecedents. Leibniz postulated a pre-established harmony for his windowless monads which mirror the same world, and the occasionalists postulated God as the coordinator of mind and body. Reichenbach was, to the best of my knowledge, the first to give a precise characterization of the conjunctive fork, and to formulate the general principle that conjunctive forks are open only to the future, not to the past.[28] The result is that we cannot explain coincidences on the basis of future effects, but only on the basis of antecedent causes. A widespread blackout is explained by a power failure, not by the looting which occurs as a consequence. (A common effect E may form a conjunctive fork with A and B, but only if there is also a common cause C.) The principle that conjunctive forks are not open to the past accounts for Russell's principle that symmetrical patterns emanate from a central source--they do not converge from afar upon the central point. It is also closely related to the operation of the second law of thermodynamics and the increase of entropy in the physical world.

The common cause principle has, I believe, deep explanatory significance. Bas van Fraassen has recently subjected it to careful scrutiny, and he has convinced me that Reichenbach's formulation in terms of the conjunctive fork, as he defined it, is faulty.[29] (We do not, however, agree about the nature of the flaw.) There are, it seems, certain sorts of causal *interactions* in which the resulting effects are more strongly correlated with one another than is allowed in Reichenbach's conjunctive forks. If, for example, an energetic photon collides with an electron in a Compton scattering experiment, there is a certain probability that a photon with a given smaller energy will emerge, and there is a certain probability that the electron will be kicked out with a given kinetic energy (see Figure 1). However, because of the law of conservation of energy, there is a strong correspondence between the two energies--their sum must be close to the energy of the incident photon. Thus, the probability of getting a photon with energy E_1 and an electron with energy E_2, where $E_1 + E_2$ is approximately equal to E (the energy of the incident photon), is much greater than the product of the probabilities of each energy occurring separately. Assume, for example, that there is a probability of 0.1 that a photon of energy E_1 will emerge if a photon of energy E impinges on a given target,

COMPTON SCATTERING

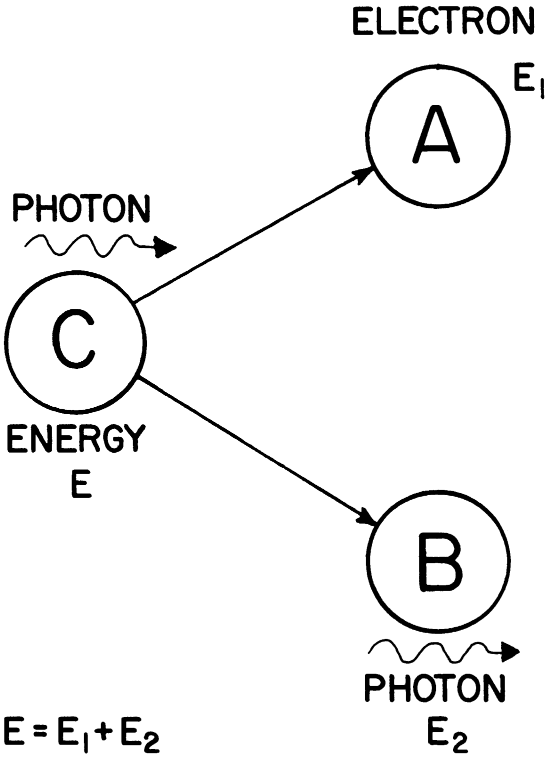

INTERACTIVE
FORK

$$\boxed{\begin{array}{c} P(A \,\&\, B/C) > \\ P(A/C) \times P(B/C) \end{array}}$$

NON-SCREENING-OFF

and assume that there is a probability of 0.1 that an electron with kinetic energy E_2 will emerge under the same circumstances (where E, E_1, and E_2 are related as the law of conservation of energy demands). In this case the probability of the joint result is not 0.01, the product of the separate probabilities, but 0.1, for each result will occur if and only if the other does.[30] The same relationships could be illustrated by such macroscopic events as collisions of billiard balls, but I have chosen Compton scattering because there is good reason to believe that events of that type are irreducibly statistical. Given a high energy photon impinging upon the electron in a given atom, there is no way, even in principle, of predicting with certainty the energies of the photon and electron which result from the interaction.

This sort of interaction stands in sharp contrast with the sort of statistical dependency we have in the leukemia example (see Figure 2, which also represents the relationships in the colorblindness case). In the absence of a strong source of radiation, such as the atomic blast, we may assume that the probability of next-door neighbors contracting the disease equals the product of the probabilities for each of them separately. If, however, we consider two next-door neighbors who lived at a distance of 2 kilometers from the hypocenter of the atomic explosion, the probability of both of them contracting leukemia is much greater than it would be for any two randomly selected members of the population at large. This apparent dependency between the two leukemia cases is not a direct physical dependency between them; it is merely a statistical result of the fact that the probability for each of them has been enhanced independently of the other by being located in close proximity to the atomic explosion. But the individual photons of radiation which impinge upon the two victims are emitted independently, travel independently, and damage living tissues independently.

It thus appears that there are two kinds of causal forks: (1) Reichenbach's *conjunctive forks*, in which the common cause screens-off the one effect from the other, which are exemplified by the colorblindness and leukemia cases, and (2) *interactive forks*, exemplified by the Compton scattering of a photon and an electron. In forks of the interactive sort, the common cause does not screen-off the one effect from the other. The probability that the electron will be ejected with kinetic energy E_2 given an incident photon of energy E is *not equal to* the probability that the electron will emerge with energy E_2 given an incident photon of energy E and a scattered photon of energy E_1. In the conjunctive fork, the common cause C absorbs the dependency between the effects A and B, for the probability of A and B given C is *equal to* the product of the probability of A given C and the probability of B given C. In the interactive fork, the common cause C does not absorb the dependency between the effects A and B, for the probability of A and B given C is *greater than* the product of the two separate conditional probabilities.[31]

Recognition and characterization of the interactive fork enables us to fill a serious lacuna in the treatment up to this point. I have discussed causal processes, indicating roughly how they are to be characterized, and I have mentioned causal interactions, but have said nothing about their characterization.

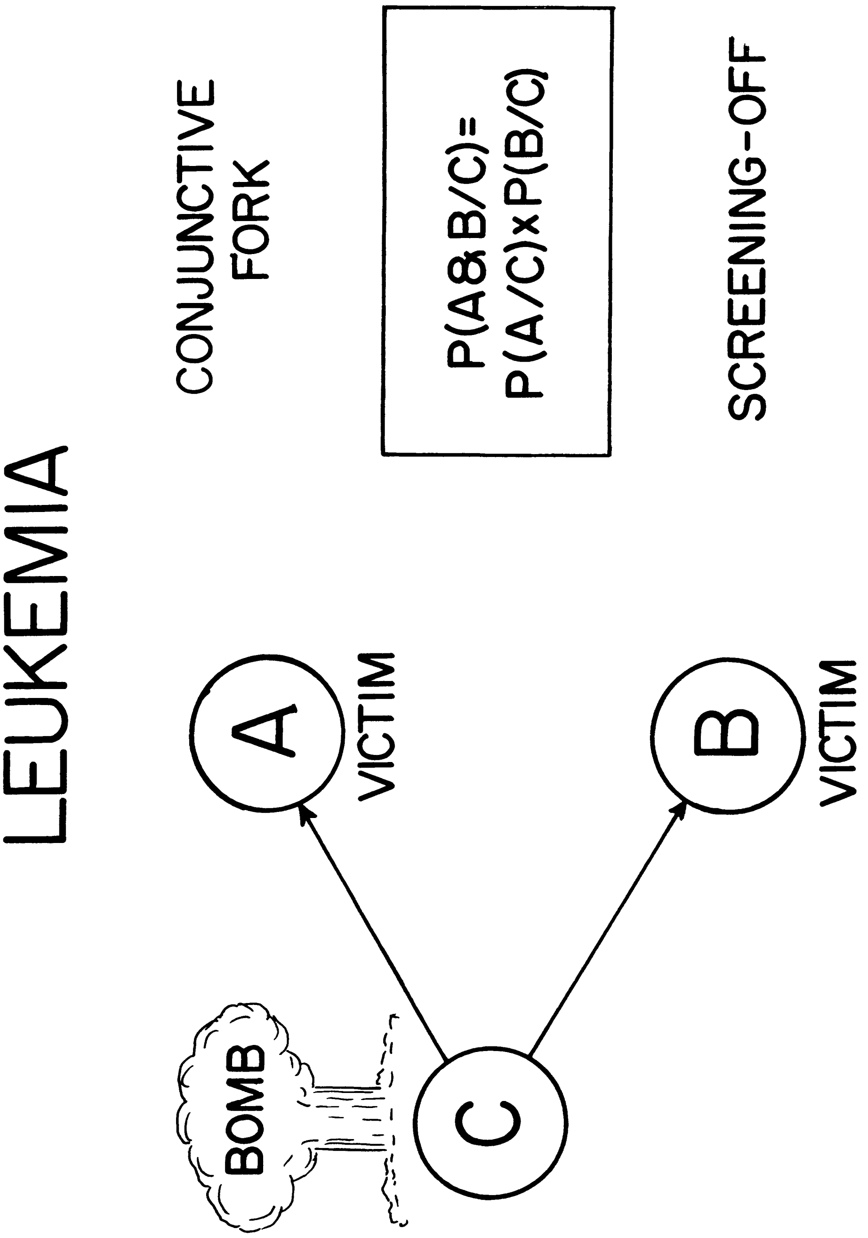
LEUKEMIA
CONJUNCTIVE FORK
P(A&B/C)=
P(A/C)×P(B/C)
SCREENING-OFF
A
VICTIM
B
VICTIM
BOMB
C

Indeed, the criterion by which we distinguished causal processes from pseudo-processes involved the use of marks, and marks are obviously results of causal interactions. Thus, our account stands in serious need of a characterization of causal interactions, and the interactive fork enables us, I believe, to furnish it.

There is a strong temptation to think of events as basic types of entities, and to construe processes--real or pseudo--as collections of events. This viewpoint may be due, at least in part, to the fact that the space-time interval between events is a fundamental invariant of the special theory of relativity, and that events thus enjoy an especially fundamental status. I suggest, nevertheless, that we reverse the approach. Let us begin with processes (which have not yet been sorted out into causal and pseudo) and look at their intersections. We can be reassured about the legitimacy of this new orientation by the fact that the basic space-time structure of both special relativity and general relativity can be built upon processes without direct recourse to events.[32] An electron traveling through space is a process, and so is a photon; if they collide, that is an intersection. A light pulse traveling from a beacon to a screen is a process, and a piece of red glass standing in the path is another; the light passing through the glass is an intersection. Both of these intersections constitute interactions. If two light beams cross one another, we have an intersection without an interaction--except in the extremely unlikely event of a particle-like collision between photons. What we want to say, very roughtly, is that when two processes intersect, and both are modified in such ways that the changes in one are correlated with changes in the other--in the manner of an interactive fork (see Figure 3)--we have a causal interaction. There are technical details to be worked out before we can claim to have a satisfactory account, but the general idea seems clear enough.[33]

I should like to commend the principle of the common cause--so construed as to make reference to both conjunctive forks and interactive forks--to your serious consideration.[34] Several of its uses have already been mentioned and illustrated. First, it supplies a schema for the straightforward explanations of everyday sorts of otherwise improbable coincidences. Second, it is the source of the fundamental temporal asymmetry of causality, and it accounts for the temporal asymmetry we impose upon scientific explanations. Third, it provides the key to the explication of the concept of causal interaction. These considerations certainly testify to its philosophical importance.

There are, however, two additional applications to which I should like to call attention. Fourth, as Russell showed, the principle plays a fundamental role in the causal theory of perception. When various observers (including cameras as well as human beings) arranged around a central region, such as a stage in theatre-in-the-round, have perceptions which correspond systematically with one another in the customary way, we may infer, with reasonable reliability, that they have a common cause--namely, a drama being performed on the stage.[35] This fact has considerable epistemological import.

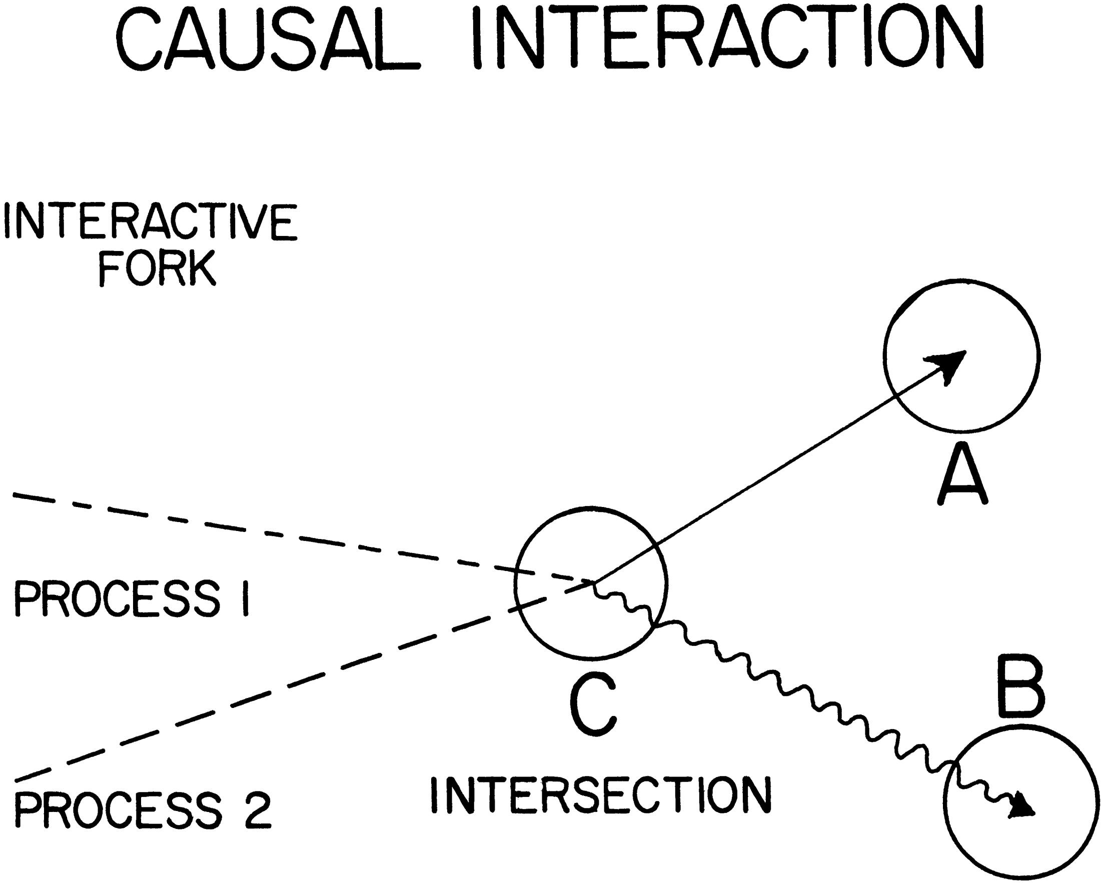
CAUSAL INTERACTION
INTERACTIVE FORK
A
PROCESS 1
C
B
PROCESS 2
INTERSECTION

Fifth, the principle of the common cause can be invoked to support scientific realism.[36] Suppose, going back to a previous example, we have postulated the existence of molecules to provide a causal explanation of the phenomena governed by the ideal gas law. We will naturally be curious about their properties--how large they are, how massive they are, how many there are. An appeal to Brownian motion enables us to infer such things. By microscopic examination of smoke particles suspended in a gas, we can ascertain their average kinetic energies, and since the observed system can be assumed to be in a state of thermal equilibrium, we can immediately infer the average kinetic energies of the molecules of the gas in which the particles are suspended. Since average velocities of the molecules are straightforwardly ascertainable by experiment, we can easily find the masses of the individual molecules, and hence, the number of molecules in a given sample of gas. If the sample consists of precisely one mole (gram molecular weight) of the particular gas, the number of molecules in the sample is Avogadro's number--a fundamental physical constant. Thus, the causal explanation of Brownian motion yields detailed quantitative information about the micro-entities of which the gas is composed.

Now, consider another phenomenon which appears to be of an altogether different sort. If an electric current is passed through an electrolytic solution--for example, one containing a silver salt--a certain amount of metallic silver is deposited on the cathode. The amount deposited is proportional to the amount of electric charge which passes through the solution. In constructing a causal explanation of this phenomenon (known as electrolysis), we postulate that charged ions travel through the solution, and that the amount of charge required to deposit a singly charged ion is equal to the charge on the electron. The magnitude of the electron charge was empirically determined through the work of J. J. Thomson and Robert Millikan. The amount of electric charge required to deposit one mole of a monovalent metal is known as the Faraday, and by experimental determination, it is equal to 96,487 coulombs. When this number is divided by the charge on the electron (-1.602×10^{-19} coulombs), the result is Avogadro's number. Indeed, the Faraday is simply Avogadro's number of electron charges.

The fundamental fact to which I wish to call attention is that the value of Avogadro's number ascertained from the analysis of Brownian motion agrees, within the limits of experimental error, with the value obtained by electrolytic measurement. Without a common causal antecedent, such agreement would constitute a remarkable coincidence. The point may be put in this way. From the molecular kinetic theory of gases we can derive the statement form, "The number of molecules in a mole of gas is _______________." From the electrochemical theory of electrolysis, we can derive the statement form, "The number of electron charges in a Faraday is ____________." The astonishing fact is that the same number fills both blanks. In my opinion, the instrumentalist cannot, with impunity, ignore what must be an amazing correspondence between what happens when one scientist is watching smoke particles dancing in a container of gas while another scientist in a different laboratory is observing the electroplating of silver. Without an underlying causal mechanism--of the sort involved in the postulation of atoms, molecules, and ions--the coincidence would be as miraculous as if the number of grapes harvested in California in any given year were equal, up to the limits of observational error, to the number of coffee beans produced in Brazil in the same year. Avogadro's number, I must add, can be ascertained in a variety of other ways as well--e.g., X-ray diffraction from crystals--which also appear to be entirely different unless we postulate the existence of atoms, molecules, and ions. The principle of the common cause thus seems to apply directly to the explanation of observable regularities by appeal to unobservable entities. In this instance, to be sure, the common cause is not some sort of event; it is rather a common constant underlying structure which manifests itself in a variety of different situations.

Let me now summarize the picture of scientific explanation I have tried to outline. If we wish to explain a particular event, such as death by leukemia of GI Joe, we begin by assembling the factors statistically relevant to that occurrence--for example, his distance from the atomic explosion, the magnitude of the blast, and the type of shelter he was in. There will be many others, no doubt, but these will do for purposes of illustration. We must also obtain the probability values associated with the relevancy relations. The statistical relevance relations are statistical regularities, and we proceed to explain them. Although this differs substantially from things I have said previously, I no longer believe that the assemblage of relevant factors provides a complete explanation--or much of anything in the way of an explanation.[37] We do, I believe, have a bona fide explanation of an event if we have a complete set of statistically relevant factors, the pertinent probability values, _and_ causal explanations of the relevance relations. Subsumption of a particular occurrence under statistical regularities--which, we recall, does not imply anything about the construction of deductive or inductive arguments--is a necessary part of any adequate explanation of its occurrence, but it is not the whole story. The causal explanation of the regularity is also needed. This claim, it should be noted, is in direct conflict with the received view, according to which the mere subsumption--deductive or inductive--of an event under a lawful regularity constitutes a complete explanation. One can, according to the received view, go on to ask for an explanation of any law used to explain a given event, but that is a different explanation. I am suggesting, on the contrary, that if the regularity invoked is not a causal regularity, then a causal explanation of that very regularity must be made part of the explanation of the event.

If we have events of two types, _A_ and _B_, whose respective members are not spatio-temporally contiguous, but whose occurrences are correlated with one another, the causal explanation of this regularity may take either of two forms. Either there is a direct causal connection from _A_ to _B_ or from _B_ to _A_, or there is a common cause _C_ which accounts for the statistical dependency. In either case, those events which stand in the cause-effect relation to one another are joined by a causal process.[38] The distinct events _A_, _B_, and _C_ which are thus related constitute interactions--as defined in terms of an interactive fork--at the appropriate places in the respective causal processes. The interactions produce modifications in the causal processes, and the causal processes transmit the modifications. Statistical dependency relations arise out of local interactions--there is no action-at-a-distance (as far as macro-phenomena are concerned, at least)--and they are propagated through the world by causal processes. In our leukemia example, a slow neutron, impinging upon a uranium atom, has a certain probability of inducing nuclear fission, and if fission occurs, gamma radiation is emitted. The gamma ray travels through space, and it may interact with a human cell, producing a modification which may leave the cell open to attack by the virus associated with leukemia. The fact that many such interactions of neutrons with fissionable nuclei are occurring in close spatio-temporal proximity, giving rise to processes which radiate in all directions, produces a pattern of statistical dependency relations. After initiation, these processes go on independently of one another, but they do produce relationships which can be described by means of the conjunctive fork.

Causal processes and causal interactions are, of course, governed by various laws--e.g., conservation of energy and momentum. In a causal process, such as the propagation of a light wave or the free motion of a material particle, energy is being transmitted. The distinction between causal processes and pseudo-processes lies in the distinction between the transmission of energy from one space-time locale to another and the mere appearance of energy at various space-time locations. When causal interactions occur--not merely intersections of processes--we have energy and/or momentum transfer. Such laws as conservation of energy and momentum are causal laws in the sense that they are regularities exhibited by causal processes and interactions.

Near the beginning, I suggested that deduction of a restricted law from a more general law constitutes a paradigm of a certain type of explanation. No theory of scientific explanation can hope to be successful unless it can handle cases of this sort. Lenz's law, for example, which governs the direction of flow of an electric current generated by a changing magnetic field, can be deduced from the law of conservation of energy. But this deductive relation shows that the more restricted regularity is simply part of a more comprehensive physical pattern expressed by the law of conservation of energy. Similarly, Kepler's laws of planetary motion describe a restricted subclass of the class of all motions governed by Newtonian mechanics. The deductive relations exhibit what amounts to a part-whole relationship, but it is, in my opinion, the physical relationship between the more comprehensive physical regularity and the less comprehensive physical regularity which has explanatory significance. I should like to put it this way. An explanation may sometimes provide the materials out of which an argument, deductive or inductive, can be constructed; an argument may sometimes exhibit explanatory relations. It does not follow, however, that explanations are arguments.

Earlier in this discussion, I mentioned three shortcomings in the most widely held theories of scientific explanation. I should now like to indicate the ways in which the theory I have been outlining attempts to cope with these problems. (1) The causal conception, I claimed, has lacked an adequate analysis of causation. The foregoing explications of causal processes and causal interactions were intended to fill that gap. (2) The inferential conception, I claimed, had misconstrued the relation of subsumption under law. When we see how statistical relevance relations can be brought to bear upon facts-to-be-explained, we discover that it is possible to have a *covering law conception* of scientific explanation without regarding explanations as arguments. The recognition that subsumption of narrower regularities under broader regularities can be viewed as a part-whole relation reinforces that point. At the same time, it suggests a reason for the tremendous appeal of the inferential conception in the first place. (3) Both of the popular conceptions, I claimed, overlooked a fundamental explanatory principle. That principle, obviously, is the principle of the common cause. I have tried to display its enormous explanatory significance. The theory outlined above is designed to overcome all three of these difficulties.

On the basis of the foregoing characterization of scientific explanation, how should we answer the question posed at the outset? What does Laplace's demon lack, if anything, with respect to the explanatory aim of science? Several items may be mentioned. The demon *may* lack an adequate recognition of the distinction between causal laws and non-causal regularities; it *may* lack adequate knowledge of causal processes and of their ability to propagate causal influence; and it *may* lack adequate appreciation of the role of causal interactions in *producing* changes and regularities in the world. None of these capabilities was explicitly demanded by Laplace, for his analysis of causal relations in general was rather superficial.

What does scientific explanation offer, over and above the inferential capacity of prediction and retrodiction, at which the Laplacian demon excelled? It provides knowledge of the mechanisms of *production* and *propagation* of structure in the world. That goes some distance beyond mere recognition of regularities, and of the possibility of subsuming particular phenomena thereunder. It is my view that knowledge of the mechanisms of production and propagation of structure in the world yields scientific understanding, and that this is what we seek when we pose explanation-seeking why questions. The answers are well worth having. That is why we ask, not only "What?" but "Why?"

NOTES

1. The author wishes to express his gratitude to the National Science Foundation for support of research on scientific explanation.

2. Book I.2, 71b, 17-24.

3. Antoine Arnauld, The Art of Thinking (Indianapolis: Bobbs-Merrill, 1964), p. 330. "Such demonstrations may convince the mind, but they do not enlighten it; and enlightenment ought to be the principal fruit of true knowledge. Our minds are unsatisfied unless they know not only that a thing is but why it is."

4. Karl Pearson, The Grammar of Science, 3rd ed. (New York: Meridian Books, 1957), p. xi. The first edition appeared in 1892, the second in 1899, and the third was first published in 1911. In the Preface to the Third Edition, Pearson remarked, just before the statement quoted in the text, "Reading the book again after many years, it was surprising to find how the heterodoxy of the 'eighties had become the commonplace and accepted doctrine of to-day." Since the "commonplace and accepted doctrine" of 1911 has again become heterodox, one wonders to what extent such changes in philosophic doctrine are mere matters of changing fashion.

5. Bertrand Russell, Our Knowledge of the External World (London: George Allen & Unwin Ltd, 1922), p. 179.

6. The classic paper by Carl G. Hempel and Paul Oppenheim, "Studies in the Logic of Explanation," which has served as the point of departure for almost all subsequent discussion was first published just thirty years ago in 1948 in Philosophy of Science, Vol. 15, pp. 135-175.

7. See, for example, his recent paper, "Causation as Explanation," Nous, Vol. 9 (1975), pp. 3-16.

8. Hempel's conceptions have been most thoroughly elaborated in his monographic essay, "Aspects of Scientific Explanation," in Aspects of Scientific Explanation and Other Essays in the Philosophy of Science (New York: Free Press, 1965), pp. 331-496.

9. P. S. Laplace, A Philosophical Essay on Probabilities (New York: Dover Publications, 1951), pp. 3-6 .

10. In "A Third Dogma of Empiricism" in Robert Butts and Jaakko Hintikka, eds., Basic Problems in Methodology and Linguistics (Dordrecht: D. Reidel Publishing Co., 1977), pp. 149-166, I have given an extended systematic critique of the thesis (dogma?) that scientific explanations are arguments.

11. Hempel, “Aspects of Scientific Explanation,” pp. 352-354.

12. Patrick Suppes, *A Probabilistic Theory of Causation* (Amsterdam: North-Holland Publishing Co., 1970); Hans Reichenbach, *The Direction of Time* (Berkeley & Los Angeles: University of California Press, 1956), Chap. IV.

13. See *Nature*, Vol. 271 (2 Feb. 1978), p. 399.

14. Irving Copi, *Introduction to Logic*, 4th ed. (New York: Macmillan Publishing Co., 1972), pp. 396-397, cites this example from *No More War* by Linus Pauling.

15. According to the article in *Nature* (note 13), “the eight reported cases of leukaemia among 2235 [soldiers] was ‘out of the normal range’.” Dr. Karl Z. Morgan “had ‘no doubt whatever’ that [the] radiation had caused the leukaemia now found in those who had taken part in the manoeuvers.”

16. Hempel’s inductive-statistical model, as formulated in “Aspects of Scientific Explanation” (1965) embodied such a high probability requirement, but in “Nachwort 1976” inserted into a German translation of this article (*Aspekte wissenschaftlicher Erklarung*, Walter de Gruyter, 1977) this requirement is retracted.

17. Bertrand Russell, *Human Knowledge, Its Scope and Limits* (New York: Simon and Schuster, 1948), p. 459.

18. Hans Reichenbach, *The Philosophy of Space and Time* (New York: Dover Publications, 1958), Sec. 21.

19. See my “Theoretical Explanation‘” Sec. 3, pp. 129-134, in Stephan Körner, ed., *Explanation* (Oxford: Basil Blackwell, 1975), for a more detailed discussion of this distinction. It is an unfortunate lacuna in Russell’s discussion of causal lines--though one which can easily be repaired--that he does not notice the distinction between causal processes and pseudo-processes.

20. See Wesley C. Salmon, ed., *Zeno’s Paradoxes* (Indianapolis: Bobbs-Merrill, 1970), p. 23, for a description of this “theory.”

21. I have made an attempt to elaborate this idea in “a ‘At-At’ Theory of Causal Influence,” *Philosophy of Science*, Vol. 44, No. 2 (June 1977), pp. 215-224. Because of a criticism due to Nancy Cartwright, I now realize that the formulation given in this article is not entirely satisfactory, but I think the difficulty can be repaired.

22. Russell, *Our Knowledge of the External World*, Lecture VI, “The Problem of Infinity Considered Historically.” The relevant portions are reprinted in my anthology, *Zeno’s Paradoxes.*

23. Russell, *Human Knowledge*, pp. 460-475.

24. Reichenbach, *The Direction of Time*, Sec. 19.

25. In "Theoretical Explanation" I discuss the explanatory import of the common cause principle in greater detail.

26. Reichenbach offers the following formal definition of a conjunctive fork ACB

$$P(A\&B/C) = P(A/C) \times P(B/C)$$

$$P(A\&B/\bar{C}) = P(A/\bar{C}) \times P(B/\bar{C})$$

$$P(A/C) > P(A/\bar{C})$$

$$P(B/C) > P(B/\bar{C})$$

in *The Direction of Time*, p. 159. I have changed these formulas from Reichenbach's notation into a more standard one.

27. C screens-off A from B if

$$P(A/C\&B) = P(A/C) \neq P(A/B)$$

28. *The Direction of Time*, pp. 162-163.

29. Bas C. van Fraassen, "The Pragmatics of Explanation," *American Philosophical Quarterly*, Vol. 14, No. 2 (April 1977), pp. 143-150. This paper was presented at the 51st Annual Meeting of the American Philosophical Association, Pacific Division, March 1977.

30. The relation between $E_1 + E_2$ and E is an approximate rather than a precise equality because the ejected electron has some energy of its own before scattering, but this energy is so small compared with the energy of the incident X-ray or γ-ray photon that it can be neglected. When I refer to the probability that the scattered photon and electron will have energies E_1 and E_2 respectively, this should be taken to mean that these energies fall within some specified interval, not that they have exact values.

31. As the boxed formulas in Figures 1 and 2 indicate, the difference between a conjunctive fork and an interactive fork lies in the difference between

$$P(A\&B/C) = P(A/C) \times P(B/C)$$

and

$$P(A\&B/C) > P(A/C) \times P(B/C).$$

The remaining formulas given in Note 26 may be incorporated into the definitions of both kinds of forks.

One reason why Reichenbach may have failed to notice the interactive fork is that, in the special case in which

$$\underline{P}(\underline{A}/\underline{C}) = \underline{P}(\underline{B}/\underline{C}) = 1,$$

the conjunctive fork shares a fundamental property of the interactive fork, namely, a perfect correlation between A and B given C. Many of his illustrative examples are instances of this special case.

32. For the special theory of relativity, this has been shown by John Winnie in "The Causal Theory of Space-time" in John S. Earman, Clark N. Glymour, and John J. Stachel, eds., Foundations of Space-Time Theories, Minnesota Studies in the Philosophy of Science, Vol. VIII (Minneapolis University of Minnesota Press, 1977) pp. 134-205, which utilizes much earlier results of A. A. Robb. For general relativity, the approach is discussed under the heading "The Geodesic Method" in Adolf Grünbaum, Philosophical Problems of Space and Time, 2nd ed. (Dordrecht: D. Reidel Publishing Co., 1973), pp. 735-750.

33. The whole idea of characterizing causal interactions in terms of forks was suggested by Philip von Bretzel in "Concerning a Probabilistic Theory of Causation Adequate for the Causal Theory of Time," Synthese, Vol. 35, No. 2 (June 1977), pp. 173-190, especially Note 13.

34. It strikes me as an unfortunate fact that this important principle seems to have gone largely unnoticed by philosophers ever since its publication in Reichenbach's The Direction of Time in 1956.

35. Russell, Human Knowledge, pp. 491-492.

36. Scientific realism is a popular doctrine nowadays, and most contemporary philosophers of science probably do not feel any pressing need for additional arguments to support this view. Although I am thoroughly convinced (in my heart) that scientific realism is correct, I am largely dissatisfied with the arguments usually brought in support of it. The argument I am about to outline seems to me more satisfactory than others.

37. Compare Wesley C. Salmon, et al., Statistical Explanation and Statistical Relevance (Pittsburgh: University of Pittsburgh Press, 1971), p. 78. There I ask, "What more could one ask of an explanation?" The present paper attempts to present at least part of the answer.

38. Reichenbach believed that various causal relations, including conjunctive forks, could be explicated entirely in terms of the statistical relations among the events involved. I do not believe this is possible; it seems to me that we must also establish the appropriate connections via causal processes.