

Philosophical Papers

Volume II

DAVID LEWIS



New York Oxford
OXFORD UNIVERSITY PRESS
1986

(2)

Oxford University Press
Oxford New York Toronto
Delhi Bombay Calcutta Madras Karachi
Petaling Jaya Singapore Hong Kong Tokyo
Nairobi Dar es Salaam Cape Town
Melbourne Auckland

and associated companies in
Beirut Berlin Ibadan Nicosia

Copyright © 1986 David Lewis

Published by Oxford University Press, Inc.,
200 Madison Avenue, New York, New York 10016

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means
electronic, mechanical, photocopying, recording, or otherwise,
without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data
(Revised for vol. 2)

Lewis, David K.
Philosophical papers.

Includes bibliographies and indexes.

1. Philosophy. I. Title.
B29.L49 1983 191 82-22551
ISBN 0-19-503203-9 (v. 1)
ISBN 0-19-503204-7 (pbk. : v. 1)
ISBN 0-19-503645-X (v. 2)
ISBN 0-19-503646-8 (pbk. : v. 2)

2 4 6 8 9 7 5 3 1

Printed in the United States of America

Acknowledgments

I thank the editors and publishers who have kindly granted permission to reprint eleven of the essays that appears in this volume. The places of first publication are as follows.

“Counterfactuals and Comparative Possibility,” *Journal of Philosophical Logic* 2 (1973): 418–46.

“Counterfactual Dependence and Time’s Arrow,” *Noûs* 13 (1979): 455–76.

“The Paradoxes of Time Travel,” *American Philosophical Quarterly* 13 (1976): 145–52.

“A Subjectivist’s Guide to Objective Chance,” in Richard C. Jeffrey, ed., *Studies in Inductive Logic and Probability*, Volume II (University of California Press, 1980).

“Probabilities of Conditionals and Conditional Probabilities,” *Philosophical Review* 85 (1976): 297–315.

“Causation,” *Journal of Philosophy* 70 (1973): 556–67; the post-scripts incorporate parts of those that are translated as “Kausalität: Epilog 1978” in Günter Posch, ed., *Kausalität—Neue Texte* (Philip Reclam, 1981).

“Veridical Hallucination and Prosthetic Vision,” *Australasian Journal of Philosophy* 58 (1980): 239–49.

“Are We Free to Break the Laws?” *Theoria* 47 (1981): 113–21.

“Prisoners’ Dilemma is a Newcomb Problem,” *Philosophy and Public Affairs* 8 (1979): 235–40.

· TWENTY-TWO ·

Causal Explanation*

I. CAUSAL HISTORIES

Any particular event that we might wish to explain stands at the end of a long and complicated causal history. We might imagine a world where causal histories are short and simple; but in the world as we know it, the only question is whether they are infinite or merely enormous.

An explanandum event has its causes. These act jointly. We have the icy road, the bald tire, the drunk driver, the blind corner, the approaching car, and more. Together, these cause the crash. Jointly they suffice to make the crash inevitable, or at least highly probable, or at least much more probable than it would otherwise have been. And the crash depends on each. Without any one it would not have happened, or at least it would have been very much less probable than it was.

But these are by no means all the causes of the crash. For one thing, each of these causes in turn has its causes; and those too are causes of the crash. So in turn are their causes, and so, perhaps, *ad infinitum*. The crash is the culmination of countless distinct, converging causal chains.

* This paper is descended, distantly, from my Hågerstrom Lectures in Uppsala in 1977, and more directly from my Howison Lectures in Berkeley in 1979.

Roughly speaking, a causal history has the structure of a tree. But not quite: the chains may diverge as well as converge. The roots in childhood of our driver's reckless disposition, for example, are part of the causal chains via his drunkenness, and also are part of other chains via his bald tire.

Further, causal chains are dense. (Not necessarily, perhaps—time might be discrete—but in the world as we mostly believe it to be.) A causal chain may go back as far as it can go and still not be complete, since it may leave out intermediate links. The blind corner and the oncoming car were not immediate causes of the crash. They caused a swerve; that and the bald tire and icy road caused a skid; that and the driver's drunkenness caused him to apply the brake, which only made matters worse And still we have mentioned only a few of the most salient stages in the last second of the causal history of the crash. The causal process was in fact a continuous one.

Finally, several causes may be lumped together into one big cause. Or one cause may be divisible into parts. Some of these parts may themselves be causes of the explanandum event, or of parts of it. (Indeed, some parts of the explanandum event itself may be causes of others.) The baldness of the tire consists of the baldness of the inner half plus the baldness of the outer half; the driver's drunkenness consists of many different disabilities, of which several may have contributed in different ways to the crash. There is no one right way—though there may be more or less natural ways—of carving up a causal history.

The multiplicity of causes and the complexity of causal histories are obscured when we speak, as we sometimes do, of *the* cause of something. That suggests that there is only one. But in fact it is commonplace to speak of "the X" when we know that there are many X's, and even many X's in our domain of discourse, as witness McCawley's sentence "the dog got in a fight with another dog." If someone says that the bald tire was the cause of the crash, another says that the driver's drunkenness was the cause, and still another says that the cause was the bad upbringing which made him so reckless, I do not think any of them disagree with me when I say that the causal history includes all three. They disagree only about which part of the causal history is most salient for the purposes of some particular inquiry. They may be looking for the most remarkable part, the most remediable or blameworthy part, the least obvious of the discoverable parts, Some parts will be salient in some contexts, others in others. Some will not be at all salient in any likely context, but they

belong to the causal history all the same: the availability of petrol, the birth of the driver's paternal grandmother, the building of the fatal road, the position and velocity of the car a split second before the impact.¹

(It is sometimes thought that only an aggregate of conditions inclusive enough to be sufficient all by itself—Mill's "whole cause"—deserves to be called "the cause." But even on this eccentric usage, we still have many deserving candidates for the title. For if we have a whole cause at one time, then also we have other whole causes at later times, and perhaps at earlier times as well.)

A causal history is a relational structure. Its *relata* are *events*: local matters of particular fact, of the sorts that may cause or be caused. I have in mind events in the most ordinary sense of the word: flashes, battles, conversations, impacts, strolls, deaths, touchdowns, falls, kisses, But also I mean to include events in a broader sense: a moving object's continuing to move, the retention of a trace, the presence of copper in a sample. (See my "Events," in this volume.)

These events may stand in various relations, for instance spatiotemporal relations and relations of part to whole. But it is their causal relations that make a causal history. In particular, I am concerned with relations of causal dependence. An event depends on others, which depend in turn on yet others, . . . ; and the events to which an event is thus linked, either directly or stepwise, I take to be its causes. Given the full structure of causal dependence, all other causal relations are given. Further, I take causal dependence itself to be counterfactual dependence, of a suitably non-backtracking sort, between distinct events: in Hume's words, "if the first . . . had not been, the second never had existed."² (See "Causation," in this volume.) But this paper is not meant to rely on my views about the analysis of causation.

¹ On definite descriptions that do not imply uniqueness, see "Scorekeeping in a Language Game," in my *Philosophical Papers*, Volume I; and James McCawley, "Pre-supposition and Discourse Structure," in *Syntax and Semantics* 11, ed. by David Dineen and Choon-kyu Oh (New York: Academic Press, 1979). On causal selection, see Morton G. White, *Foundations of Historical Knowledge* (New York: Harper & Row, 1965), Chapter IV. Peter Unger, in "The Uniqueness of Causation," *American Philosophical Quarterly* 14 (1977): 177–88, has noted that not only "the cause of" but also the verb "caused" may be used selectively. There is something odd—inconsistent, he thinks—in saying with emphasis that each of two distinct things caused something. Even "a cause of" may carry some hint of selectivity. It would be strange, though I think not false, to say in any ordinary context that the availability of petrol was a cause of the crash.

² *An Enquiry Concerning Human Understanding*, Section VII.

Whatever causation may be, there are still causal histories, and what I shall say about causal explanation should still apply.³

I include relations of probabilistic causal dependence. Those who know of the strong scientific case for saying that our world is an indeterministic one, and that most events therein are to some extent matters of chance, never seriously renounce the commonsensical view that there is plenty of causation in the world. (They may preach the "downfall of causality" in their philosophical moments. But whatever that may mean, evidently it does not imply any shortage of causation.) For instance, they would never dream of agreeing with those ignorant tribes who disbelieve that pregnancies are caused by events of sexual intercourse. The causation they believe in must be probabilistic. And if, as seems likely, our world is indeed thoroughly indeterministic and chancy, its causal histories must be largely or entirely structures of probabilistic causal dependence. I take such dependence to obtain when the objective chances of some events depend counterfactually upon other events: if the cause had not been, the effect would have been very much less probable than it actually was. (See Postscript B to "Causation," in this volume.) But again, what is said in this paper should be compatible with any analysis of probabilistic causation.

The causal history of a particular event includes that event itself, and all events which are part of it. Further, it is closed under causal dependence: anything on which an event in the history depends is itself an event in the history. (A causal history need not be closed under the converse relation. Normally plenty of omitted events will depend on included ones.) Finally, a causal history includes no more than it must to meet these conditions.

II. EXPLANATION AS INFORMATION

Here is my main thesis: *to explain an event is to provide some information about its causal history.*

In an act of explaining, someone who is in possession of some infor-

³ One author who connects explanation and causation in much the same way that I do, but builds on a very different account of causation, is Wesley C. Salmon. See his "Theoretical Explanation," in *Explanation*, ed. by Stephen Körner (New Haven: Yale University Press, 1975); "A Third Dogma of Empiricism," in *Basic Problems in Methodology and Linguistics*, ed. by R. Butts and J. Hintikka (Dordrecht: Reidel, 1977); and "Why Ask 'Why?'" *Proceedings of the American Philosophical Association* 51 (1978): 683–705.

mation about the causal history of some event—*explanatory information*, I shall call it—tries to convey it to someone else. Normally, to someone who is thought not to possess it already, but there are exceptions: examination answers and the like. Afterward, if the recipient understands and believes what he is told, he too will possess the information. The why-question concerning a particular event is a request for explanatory information, and hence a request that an act of explaining be performed.

In one sense of the word, an explanation of an event is such an act of explaining. To quote Sylvain Bromberger, “an explanation may be something about which it makes sense to ask: How long did it take? Was it interrupted at any point? Who gave it? When? Where? What were the exact words used? For whose benefit was it given?”⁴ But it is not clear whether just any act of explaining counts as an explanation. Some acts of explaining are unsatisfactory; for instance the explanatory information provided might be incorrect, or there might not be enough of it, or it might be stale news. If so, do we say that the performance was no explanation at all? Or that it was an unsatisfactory explanation? The answer, I think, is that we will gladly say either—thereby making life hard for those who want to settle, once and for all, the necessary and sufficient conditions for something to count as an explanation. Fortunately that is a project we needn’t undertake.

Bromberger goes on to say that an explanation “may be something about which none of [the previous] questions makes sense, but about which it makes sense to ask: Does anyone know it? Who thought of it first? Is it very complicated?” An explanation in this second sense of the word is not an act of explaining. It is a chunk of explanatory information—information that may once, or often, or never, have been conveyed in an act of explaining. (It might even be information that never could be conveyed, for it might have no finite expression in any language we could ever use.) It is a proposition about the causal history of the explanandum event. Again it is unclear—and again we needn’t make it clear—what to say about an unsatisfactory chunk of explanatory information, say one that is incorrect or one that is too small to suit us. We may call it a bad explanation, or no explanation at all.

Among the true propositions about the causal history of an event, one is maximal in strength. It is the whole truth on the subject—the biggest chunk of explanatory information that is free of error. We

⁴ “An Approach to Explanation,” in *Analytical Philosophy: Second Series*, ed. by R. J. Butler (Oxford: Blackwell, 1965).

might call this the *whole* explanation of the explanandum event, or simply *the* explanation. (But “the explanation” might also denote that one out of many explanations, in either sense, that is most salient in a certain context.) It is, of course, very unlikely that so much explanatory information ever could be known, or conveyed to anyone in some tremendous act of explaining!

One who explains may provide not another, but rather himself, with explanatory information. He may think up some hypothesis about the causal history of the explanandum event, which hypothesis he then accepts. Thus Holmes has explained the clues (correctly or not, as the case may be) when he has solved the crime to his satisfaction, even if he keeps his solution to himself. His achievement in this case probably could not be called “an explanation”; though the chunk of explanatory information he has provided himself might be so called, especially if it is a satisfactory one.

Not only a person, but other sorts of things as well, may explain. A theory or a hypothesis, or more generally any collection of premises, may provide explanatory information (correct or incorrect) by implying it. That is so whether or not anyone draws the inference, whether or not anyone accepts or even thinks of the theory in question, and whether or not the theory is true. Thus we may wonder whether our theories explain more than we will ever realize, or whether other undreamt-of theories explain more than the theories we accept.

Explanatory information comes in many shapes and sizes. Most simply, an explainer might give information about the causal history of the explanandum by saying that a certain particular event is included therein. That is, he might specify one of the causes of the explanandum. Or he might specify several. And if so, they might comprise all or part of a cross-section of the causal history: several events, more or less simultaneous and causally independent of one another, that jointly cause the explanandum. Alternatively, he might trace a causal chain. He might specify a sequence of events in the history, ending with the explanandum, each of which is among the causes of the next. Or he might trace a more complicated, branching structure that is likewise embedded in the complete history.

An explainer well might be unable to specify fully any particular event in the history, but might be in a position to make existential statements. He might say, for instance, that the history includes an event of such-and-such kind. Or he might say that the history includes several events of such-and-such kinds, related to one another in such-and-such ways. In other words, he might make an existential statement

to the effect that the history includes a pattern of events of a certain sort. (Such a pattern might be regarded, at least in some cases, as one complex and scattered event with smaller events as parts.) He might say that the causal history has a certain sort of cross-section, for instance, or that it includes a certain sort of causal chain.

If someone says that the causal history includes a pattern of events having such-and-such description, there are various sorts of description that he might give. A detailed structural specification might be given, listing the kinds and relations of the events that comprise the pattern. But that is not the only case. The explainer might instead say that the pattern that occupies a certain place in the causal history is some biological, as opposed to merely chemical, process. Or he might say that it has some global structural feature: it is a case of underdamped negative feedback, a dialectical triad, or a resonance phenomenon. (And he might have reason to say this even if he has no idea, for instance, what sort of thing it is that plays the role of a damper in the system in question.) Or he might say that it is a process analogous to some other, familiar process. (So in this special case, at least, there is something to the idea that we may explain by analogizing the unfamiliar to the familiar. At this point I am indebted to David Velleman.) Or he might say that the causal process, whatever it may be, is of a sort that tends in general to produce a certain kind of effect. I say "we have lungs because they keep us alive"; my point being that lungs were produced by that process, whatever it may be, that can and does produce all manner of life-sustaining organs. (In conveying that point by those words, of course I am relying on the shared presupposition that such a process exists. In explaining, as in other communication, literal meaning and background work together.) And I might say this much, whether or not I have definite opinions about what sort of process it is that produces life-sustaining organs. My statement is neutral between evolution, creation, vital forces, or what have you; it is also neutral between opinionation and agnosticism.

In short: information about what the causal history includes may range from the very specific to the very abstract. But we are still not done. There is also negative information: information about what the causal history does *not* include. "Why was the CIA man there when His Excellency dropped dead?—Just coincidence, believe it or not." Here the information given is negative, to the effect that a certain sort of pattern of events—namely, a plot—does not figure in the causal history. (At least, not in that fairly recent part where one might have been suspected. Various ancient plots doubtless figure in the causal histories of all current events, this one included.)

A final example. The patient takes opium and straightway falls asleep; the doctor explains that opium has a dormitive virtue. Doubtless the doctor's statement was not as informative as we might have wished, but observe that it is not altogether devoid of explanatory information. The test is that it suffices to rule out at least some hypotheses about the causal history of the explanandum. It rules out this one: the opium merchants know that opium is an inert substance, yet they wish to market it as a soporific. So they keep close watch; and whenever they see a patient take opium, they sneak in and administer a genuine soporific. The doctor has implied that this hypothesis, at least, is false; whatever the truth may be, at least it somehow involves distinctive intrinsic properties of the opium.

Of course I do not say that all explanatory information is of equal worth; or that all of it equally deserves the honorific name "explanation." My point is simply that we should be aware of the variety of explanatory information. We should not suppose that the only possible way to give some information about how an event was caused is to name one or more of its causes.

III. NON-CAUSAL EXPLANATION?

It seems quite safe to say that the provision of information about causal histories figures very prominently in the explaining of particular events. What is not so clear is that it is the whole story. Besides the causal explanation that I am discussing, is there also any such thing as non-causal explanation of particular events? My main thesis says there is not. I shall consider three apparent cases of it, one discussed by Hempel and two suggested to me by Peter Railton.⁵

First case. We have a block of glass of varying refractive index. A beam of light enters at point *A* and leaves at point *B*. In between, it passes through point *C*. Why? Because *C* falls on the path from *A* to *B* that takes light the least time to traverse; and according to Fermat's principle of

⁵ Carl G. Hempel, *Aspects of Scientific Explanation and other Essays in the Philosophy of Science* (New York: Free Press, 1965), p. 353; Peter Railton, *Explaining Explanation* (Ph. D. dissertation, Princeton University, 1979). I am much indebted to Railton throughout this paper, both where he and I agree and where we do not. For his own views on explanation, see also his "A Deductive-Nomological Model of Probabilistic Explanation," *Philosophy of Science* 45 (1978): 206–26; and "Probability, Explanation, and Information," *Synthese* 48 (1981): 233–56.

to the effect that the history includes a pattern of events of a certain sort. (Such a pattern might be regarded, at least in some cases, as one complex and scattered event with smaller events as parts.) He might say that the causal history has a certain sort of cross-section, for instance, or that it includes a certain sort of causal chain.

If someone says that the causal history includes a pattern of events having such-and-such description, there are various sorts of description that he might give. A detailed structural specification might be given, listing the kinds and relations of the events that comprise the pattern. But that is not the only case. The explainer might instead say that the pattern that occupies a certain place in the causal history is some biological, as opposed to merely chemical, process. Or he might say that it has some global structural feature: it is a case of underdamped negative feedback, a dialectical triad, or a resonance phenomenon. (And he might have reason to say this even if he has no idea, for instance, what sort of thing it is that plays the role of a damper in the system in question.) Or he might say that it is a process analogous to some other, familiar process. (So in this special case, at least, there is something to the idea that we may explain by analogizing the unfamiliar to the familiar. At this point I am indebted to David Velleman.) Or he might say that the causal process, whatever it may be, is of a sort that tends in general to produce a certain kind of effect. I say “we have lungs because they keep us alive”; my point being that lungs were produced by that process, whatever it may be, that can and does produce all manner of life-sustaining organs. (In conveying that point by those words, of course I am relying on the shared presupposition that such a process exists. In explaining, as in other communication, literal meaning and background work together.) And I might say this much, whether or not I have definite opinions about what sort of process it is that produces life-sustaining organs. My statement is neutral between evolution, creation, vital forces, or what have you; it is also neutral between opinionation and agnosticism.

In short: information about what the causal history includes may range from the very specific to the very abstract. But we are still not done. There is also negative information: information about what the causal history does *not* include. “Why was the CIA man there when His Excellency dropped dead?—Just coincidence, believe it or not.” Here the information given is negative, to the effect that a certain sort of pattern of events—namely, a plot—does not figure in the causal history. (At least, not in that fairly recent part where one might have been suspected. Various ancient plots doubtless figure in the causal histories of all current events, this one included.)

A final example. The patient takes opium and straightway falls asleep; the doctor explains that opium has a dormitive virtue. Doubtless the doctor’s statement was not as informative as we might have wished, but observe that it is not altogether devoid of explanatory information. The test is that it suffices to rule out at least some hypotheses about the causal history of the explanandum. It rules out this one: the opium merchants know that opium is an inert substance, yet they wish to market it as a soporific. So they keep close watch; and whenever they see a patient take opium, they sneak in and administer a genuine soporific. The doctor has implied that this hypothesis, at least, is false; whatever the truth may be, at least it somehow involves distinctive intrinsic properties of the opium.

Of course I do not say that all explanatory information is of equal worth; or that all of it equally deserves the honorific name “explanation.” My point is simply that we should be aware of the variety of explanatory information. We should not suppose that the only possible way to give some information about how an event was caused is to name one or more of its causes.

III. NON-CAUSAL EXPLANATION?

It seems quite safe to say that the provision of information about causal histories figures very prominently in the explaining of particular events. What is not so clear is that it is the whole story. Besides the causal explanation that I am discussing, is there also any such thing as non-causal explanation of particular events? My main thesis says there is not. I shall consider three apparent cases of it, one discussed by Hempel and two suggested to me by Peter Railton.⁵

First case. We have a block of glass of varying refractive index. A beam of light enters at point *A* and leaves at point *B*. In between, it passes through point *C*. Why? Because *C* falls on the path from *A* to *B* that takes light the least time to traverse; and according to Fermat’s principle of

⁵ Carl G. Hempel, *Aspects of Scientific Explanation and other Essays in the Philosophy of Science* (New York: Free Press, 1965), p. 353; Peter Railton, *Explaining Explanation* (Ph. D. dissertation, Princeton University, 1979). I am much indebted to Railton throughout this paper, both where he and I agree and where we do not. For his own views on explanation, see also his “A Deductive-Nomological Model of Probabilistic Explanation,” *Philosophy of Science* 45 (1978): 206–26; and “Probability, Explanation, and Information,” *Synthese* 48 (1981): 233–56.

least time, that is the path that any light going from *A* to *B* must follow. That seems non-causal. The light does not get to *C* because it looks ahead, calculates the path of least time to its destination *B*, and steers accordingly! The refractive index in parts of the glass that the light has not yet reached has nothing to do with causing it to get to *C*, but that is part of what makes it so that *C* is on the path of least time from *A* to *B*.

I reply that it is by no means clear that the light's passing through *C* has been explained. But if it has, that is because this explanation combines with information that its recipient already possesses to imply something about the causal history of the explanandum. Any likely recipient of an explanation that mentions Fermat's principle must already know a good deal about the propagation of light. He probably knows that the bending of the beam at any point depends causally on the local variation of refractive index around that point. He probably knows, or at least can guess, that Fermat's principle is somehow provable from some law describing that dependence together with some law relating refractive index to speed of light. Then he knows this: (1) the pattern of variation of the refractive index along some path from *A* to *C* is part of the causal history of the light's passing through *C*, and (2) the pattern is such that it, together with a pattern of variation elsewhere that is not part of the causal history, makes the path from *A* to *C* be part of a path of least time from *A* to *B*. To know this much is not to know just what the pattern that enters into the causal history looks like, but it is to know something—something relational—about that pattern. So the explanation does indeed provide a peculiar kind of information about the causal history of the explanandum, on condition that the recipient is able to supply the extra premises needed.

Second case. A star has been collapsing, but the collapse stops. Why? Because it's gone as far as it can go. Any more collapsed state would violate the Pauli Exclusion Principle. It's not that anything caused it to stop—there was no countervailing pressure, or anything like that. There was nothing to keep it out of a more collapsed state. Rather, there just was no such state for it to get into. The state-space of physical possibilities gave out. (If ordinary space had boundaries, a similar example could be given in which ordinary space gives out and something stops at the edge.)

I reply that information about the causal history of the stopping has indeed been provided, but it was information of an unexpectedly negative sort. It was the information that the stopping had no causes at all, except for all the causes of the collapse which was a precondition of the stopping. Negative information is still information. If you request

information about arctic penguins, the best information I can give you is that there aren't any.

Third case. Walt is immune to smallpox. Why? Because he possesses antibodies capable of killing off any smallpox virus that might come along. But his possession of antibodies doesn't *cause* his immunity. It is his immunity. Immunity is a disposition, to have a disposition is to have something or other that occupies a certain causal role, and in Walt's case what occupies the role is his possession of antibodies.

I reply that it's as if we'd said it this way: Walt has some property that protects him from smallpox. Why? Because he possesses antibodies, and possession of antibodies is a property that protects him from smallpox. Schematically: Why is it that something is *F*? Because *A* is *F*. An existential quantification is explained by providing an instance. I agree that something has been explained, and not by providing information about its causal history. But I don't agree that any particular event has been non-causally explained. The case is outside the scope of my thesis. That which protects Walt—namely, his possession of antibodies—is indeed a particular event. It is an element of causal histories; it causes and is caused. But that was not the explanandum. We could no more explain that just by saying that Walt possesses antibodies than we could explain an event just by saying that it took place. What we did explain was something else: the fact that something or other protects Walt. The obtaining of this existential fact is not an event. It cannot be caused. Rather, events that would provide it with a truth-making instance can be caused. We explain the existential fact by identifying the truth-making instance, by providing information about the causal history thereof, or both. (For further discussion of explanation of facts involving the existence of patterns of events, see Section VIII of "Events," in this volume.)

What more we say about the case depends on our theory of dispositions.⁶ I take for granted that a disposition requires a causal basis: one has the disposition iff one has a property that occupies a certain

⁶ See the discussions of dispositions and their bases in D. M. Armstrong, *A Materialist Theory of the Mind* (London: Routledge & Kegan Paul, 1968), pp. 85–88; Armstrong, *Belief, Truth and Knowledge* (Cambridge: Cambridge University Press, 1973), pp. 11–16; Elizabeth W. Prior, Robert Pargetter, and Frank Jackson, "Three Theses about Dispositions," *American Philosophical Quarterly* 19 (1982): 251–57; and Elizabeth W. Prior, *Dispositions* (Aberdeen: Aberdeen University Press, 1985). See also Section VIII of "Events," in this volume. Parallel issues arise for functionalist theories of mind. See my "An Argument for the Identity Theory" and "Mad Pain and Martian Pain," in *Philosophical Papers*, Volume I; and Jackson, Pargetter, and Prior, "Functionalism and Type-Type Identity Theories," *Philosophical Studies* 42 (1982): 209–25.

causal role. (I would be inclined to require that this be an intrinsic property, but that is controversial.) Shall we then identify the disposition with its basis? That would make the disposition a cause of its manifestations, since the basis is. But the identification might vary from case to case. (It surely would, if we count the unactualized cases.) For there might be different bases in different cases. Walt might be disposed to remain healthy if exposed to virus on the basis of his possession of antibodies, but Milt might be so disposed on the basis of his possession of dormant antibody-makers. Then if the disposition is the basis, immunity is different properties in the cases of Walt and Milt. Or better: "immunity" denotes different properties in the two cases, and there is no property of immunity *simpliciter* that Walt and Milt share.

That is disagreeably odd. But Walt and Milt do at least share something: the existential property of having some basis or other. This is the property such that, necessarily, it belongs to an individual *X* iff *X* has some property that occupies the appropriate role in *X*'s case. So perhaps we should distinguish the disposition from its various bases, and identify it rather with the existential property. That way, "immunity" could indeed name a property shared by Walt and Milt. But this alternative has a disagreeable oddity of its own. The existential property, unlike the various bases, is too disjunctive and too extrinsic to occupy any causal role. There is no event that is essentially a having of the existential property; *a fortiori*, no such event ever causes anything. (Compare the absurd double-counting of causes that would ensue if we said, for instance, that when a match struck in the evening lights, one of the causes of the lighting is an event that essentially involves the property of being struck in the evening or twirled in the morning. I say there is no such event.) So if the disposition is the existential property, then it is causally impotent. On this theory, we are mistaken whenever we ascribe effects to dispositions.

Fortunately we needn't decide between the two theories. Though they differ on the analysis of disposition-names like "immunity," they agree about what entities there are. There is one genuine event—Walt's possession of antibodies. There is a truth about Walt to the effect that he has the existential property. But there is no second event that is essentially a having of the existential property, but is not essentially a having of it in any particular way. Whatever "Walt's immunity" may denote, it does not denote such an event. And since there is no such event at all, there is no such event to be non-causally explained.

IV. GENERAL EXPLANATION

My main thesis concerns the explanation of particular events. As it stands, it says nothing about what it is to explain general kinds of events. However, it has a natural extension. All the events of a given kind have their causal histories, and these histories may to some extent be alike. Especially, the final parts of the histories may be much the same from one case to the next, however much the earlier parts may differ. Then information may be provided about what is common to all the parallel causal histories—call it *general explanatory information* about events of the given kind. To explain a kind of event is to provide some general explanatory information about events of that kind.

Thus explaining why struck matches light in general is not so very different from explaining why some particular struck match lit. In general, and in the particular case, the causal history involves friction, small hot spots, liberation of oxygen from a compound that decomposes when hot, local combustion of a heated inflammable substance facilitated by this extra oxygen, further heat produced by this combustion, and so on.

There are intermediate degrees of generality. If we are not prepared to say that every event of such-and-such kind, without exception, has a causal history with so-and-so features, we need not therefore abjure generality altogether and stick to explaining events one at a time. We may generalize modestly, without laying claim to universality, and say just that quite often an event of such-and-such kind has a causal history with so-and-so features. Or we may get a bit more ambitious and say that it is so in most cases, or at least in most cases that are likely to arise under the circumstances that prevail hereabouts. Such modest generality may be especially characteristic of history and the social sciences; but it appears also in the physical sciences of complex systems, such as meteorology and geology. We may be short of known laws to the effect that storms with feature *X* always do *Y*, or always have a certain definite probability of doing *Y*. Presumably there are such laws, but they are too complicated to discover either directly or by derivation from first principles. But we do have a great deal of general knowledge of the sorts of causal processes that commonly go on in storms.

The pursuit of general explanations may be very much more widespread in science than the pursuit of general laws. And not necessarily because we doubt that there are general laws to pursue. Even if the scientific community unanimously believed in the existence of power-

ful general laws that govern all the causal processes of nature, and whether or not those laws were yet known, meteorologists and geologists and physiologists and historians and engineers and laymen would still want general knowledge about the sorts of causal processes that go on in the systems they study.

V. EXPLAINING WELL AND BADLY

An act of explaining may be more or less satisfactory, in several different ways. It will be instructive to list them. It will *not* be instructive to fuss about whether an unsatisfactory act of explaining, or an unsatisfactory chunk of explanatory information, deserves to be so-called, and I shall leave all such questions unsettled.

1. An act of explaining may be unsatisfactory because the explanatory information provided is unsatisfactory. In particular, it might be misinformation: it might be a false proposition about the causal history of the explanandum. This defect admits of degree. False is false, but a false proposition may or may not be close to the truth.⁷ If it has a natural division into conjuncts, more or fewer of them may be true. If it has some especially salient consequences, more or fewer of those may be true. The world as it is may be more or less similar to the world as it would be if the falsehood were true.

2. The explanatory information provided may be correct, but there may not be very much of it. It might be a true but weak proposition; one that excludes few (with respect to some suitable measure) of the alternative possible ways the causal history of the explanandum might be. Or the information provided might be both true and strong, but unduly disjunctive. The alternative possibilities left open might be too widely scattered, too different from one another. These defects too

⁷ The analysis of verisimilitude has been much debated. A good survey is Ilkka Niiniluoto, "Truthlikeness: Comments on Recent Discussion," *Synthese* 38 (1978): 281–329. Some plausible analyses have failed disastrously, others conflict with one another. One conclusion that emerges is that it is probably a bad move to try to define a single virtue of verisimilitude-cum-strength. It's hard to say whether strength is a virtue in the case of false information, especially if we have no uniquely natural way of splitting the misinformation into true and false parts. Another conclusion is that even if this lumping together is avoided, verisimilitude still seems to consist of several distinguishable virtues.

admit of degree. Other things being equal, it is better if more correct explanatory information is provided, and it is better if that information is less disjunctive, up to the unattainable limit in which the whole explanation is provided and there is nothing true and relevant left to add.

3. The explanatory information provided may be correct, but not thanks to the explainer. He may have said what he did not know and had no very good reason to believe. If so, the act of explaining is not fully satisfactory, even if the information provided happens to be satisfactory.

4. The information provided, even if satisfactory in itself, may be stale news. It may add little or nothing to the information the recipient possesses already.

5. The information provided may not be of the sort the recipient most wants. He may be especially interested in certain parts of the causal history, or in certain questions about its overall structure. If so, no amount of explanatory information that addresses itself to the wrong questions will satisfy his wants, even if it is correct and strong and not already in his possession.

6. Explanatory information may be provided in such a way that the recipient has difficulty in assimilating it, or in disentangling the sort of information he wants from all the rest. He may be given more than he can handle, or he may be given it in a disorganized jumble.⁸ Or he may be given it in so unconvincing a way that he doesn't believe what he's told. If he is hard to convince, just telling him may not be an effective way to provide him with information. You may have to argue for what you tell him, so that he will have reason to believe you.

7. The recipient may start out with some explanatory misinformation, and the explainer may fail to set him right.

This list covers much that philosophers have said about the merits and demerits of explanations, or about what does and what doesn't deserve the name. And yet I have not been talking specifically about explanation at all! What I have been saying applies just as well to acts of providing information about *any* large and complicated structure. It might as well have been the rail and tram network of Melbourne rather than the causal history of some explanandum event. The information provided, and the act of providing it, can be satisfactory or not in pre-

⁸ As in the square peg example of Hilary Putnam, "Philosophy and our Mental Life," in his *Mind, Language and Reality* (Cambridge: Cambridge University Press, 1975), pp. 295–97.

cisely the same ways. There is no special subject: pragmatics of explanation.

Philosophers have proposed further *desiderata*. A good explanation ought to show that the explanandum event had to happen, given the laws and the circumstances; or at least that it was highly probable, and could therefore have been expected if we had known enough ahead of time; or at least that it was less surprising than it may have seemed. A good explanation ought to show that the causal processes at work are of familiar kinds; or that they are analogous to familiar processes; or that they are governed by simple and powerful laws; or that they are not too miscellaneous. But I say that a good explanation ought to show none of these things unless they are true. If one of these things is false in a given case, and if the recipient is interested in the question of whether it is true, or mistakenly thinks that it is true, then a good explanation ought to show that it is false. But that is nothing special: it falls under points 1, 5, and 7 of my list.

It is as if someone thought that a good explanation of any current event had to be one that revealed the sinister doings of the CIA. When the CIA really does play a part in the causal history, we would do well to tell him about it: we thereby provide correct explanatory information about the part of the causal history that interests him most. But in case the CIA had nothing to do with it, we ought not to tell him that it did. Rather we ought to tell him that it didn't. Telling him what he hopes to hear is not even a merit to be balanced off against the demerit of falsehood. In itself it has no merit at all. What does have merit is addressing the right question.

This much is true. We are, and we ought to be, biased in favor of believing hypotheses according to which what happens is probable, is governed by simple laws, and so forth. That is relevant to the credibility of explanatory information. But credibility is not a separate merit alongside truth; rather, it is what we go for in seeking truth as best we can.

Another proposed *desideratum* is that a good explanation ought to produce understanding. If understanding involves seeing the causal history of the explanandum as simple, familiar, or whatnot, I have already registered my objection. But understanding why an event took place might, I think, just mean possession of explanatory information about it—the more of that you possess, the better you understand. If so, of course a good explanation produces understanding. It produces possession of that which it provides. But this *desideratum*, so construed, is empty. It adds nothing to our understanding of explanation.

VI. WHY-QUESTIONS, PLAIN AND CONTRASTIVE

A why-question, I said, is a request for explanatory information. All questions are requests for information of some or other sort.⁹ But there is a distinction to be made. Every question has a maximal true answer: the whole truth about the subject matter on which information is requested, to which nothing could be added without irrelevancy or error. In some cases it is feasible to provide these maximal answers. Then we can reasonably hope for them, request them, and settle for nothing less. “Who done it?—Professor Plum.” There's no more to say.

In other cases it isn't feasible to provide maximal true answers. There's just too much true information of the requested sort to know or to tell. Then we do not hope for maximal answers and do not request them, and we always settle for less. The feasible answers do not divide sharply into complete and partial. They're all partial, but some are more partial than others. There's only a fuzzy line between enough and not enough of the requested information. “What's going on here?”—No need to mention that you're digesting your dinner. “Who is Bob Hawke?”—No need to write the definitive biography. Less will be a perfectly good answer. Why-questions, of course, are among the questions that inevitably get partial answers.

When partial answers are the order of the day, questioners have their ways of indicating how much information they want, or what sort. “In a word, what food do penguins eat?” “Why, in economic terms, is there no significant American socialist party?”

One way to indicate what sort of explanatory information is wanted is through the use of contrastive why-questions. Sometimes there is an explicit “rather than. . . .” Then what is wanted is information about the causal history of the explanandum event, not including information that would also have applied to the causal histories of alternative events, of the sorts indicated, if one of them had taken place instead. In other words, information is requested about the difference between the actualized causal history of the explanandum and the unactualized causal histories of its unactualized alternatives. Why did I visit Melbourne in 1979, rather than Oxford or Uppsala or Wellington? Because Monash University invited me. That is part of the causal

⁹ Except perhaps for questions that take imperative answers: “What do I do now, Boss?”

history of my visiting Melbourne; and if I had gone to one of the other places instead, presumably that would not have been part of the causal history of my going there. It would have been wrong to answer: Because I like going to places with good friends, good philosophy, cool weather, nice scenery, and plenty of trains. That liking is also part of the causal history of my visiting Melbourne, but it would equally have been part of the causal history of my visiting any of the other places, had I done so.

The same effect can be achieved by means of contrastive stress. Why did I *fly* to Brisbane when last I went there? I had my reasons for wanting to get there, but I won't mention those because they would have been part of the causal history no matter how I'd travelled. Instead I'll say that I had too little time to go by train. If I had gone by train, my having too little time could not have been part of the causal history of my so doing.

If we distinguish plain from contrastive why-questions, we can escape a dilemma about explanation under indeterminism. On the one hand, we seem quite prepared to offer explanations of chance events. Those of us who think that chance is all-pervasive (as well as those who suspend judgment) are no less willing than the staunchest determinist to explain the events that chance to happen.¹⁰ On the other hand, we balk at the very idea of explaining why a chance event took place—for is it not the very essence of chance that one thing happens rather than another for no reason whatsoever? Are we of two minds?

No; I think we are right to explain chance events, yet we are right also to deny that we can ever explain why a chance process yields one outcome rather than another. According to what I've already said, indeed we cannot explain why one happened *rather than the other*. (That is so regardless of the respective probabilities of the two.) The actual causal history of the actual chance outcome does not differ at all

¹⁰ A treatment of explanation in daily life, or in history, dare not set aside the explanation of chance events as a peculiarity arising only in quantum physics. If current scientific theory is to be trusted, chance events are far from exceptional. The misguided hope that determinism might prevail in history if not in physics well deserves Railton's mockery: "All but the most basic regularities of the universe stand forever in peril of being interrupted or upset by intrusion of the effects of random processes The success of a social revolution might appear to be explained by its overwhelming popular support, but this is to overlook the revolutionaries' luck: if all the naturally unstable nuclides on earth had commenced spontaneous nuclear fission in rapid succession, the triumph of the people would never have come to pass." ("A Deductive-Nomological Model of Probabilistic Explanation," pp. 223–24.) On the same point, see my Postscript B to "A Subjectivist's Guide to Objective Chance," in this volume.

from the unactualized causal history that the other outcome would have had, if that outcome had happened. A contrastive why-question with "rather" requests information about the features that differentiate the actual causal history from its counterfactual alternative. There are no such features, so the question can have no positive answer. Thus we are right to call chance events inexplicable, if it is contrastive explanation that we have in mind. (Likewise, we can never explain why a chance event *had* to happen, because it didn't have to.) But take away the "rather" (and the "had") and explanation becomes possible. Even a chance event has a causal history. There is information about that causal history to be provided in answer to a plain why-question. And thus we are right to proceed as we all do in explaining what we take to be chance events.

VII. THE COVERING-LAW MODEL

The covering-law model of explanation has long been the leading approach. As developed in the work of Hempel and others, it is an elegant and powerful theory. How much of it is compatible with what I have said?

Proponents of the covering-law model do not give a central place to the thesis that we explain by providing information about causes. But neither do they say much against it. They may complain that the ordinary notion of causation has resisted precise analysis; they may say that mere mention of a cause provides less in the way of explanation than might be wished; they may insist that there are a few special cases in which we have good non-causal explanations of particular occurrences. But when they give us their intended examples of covering-law explanation, they almost always pick examples in which—as they willingly agree—the covering-law explanation does include a list of joint causes of the explanandum event, and thereby provides information about its causal history.

The foremost version of the covering-law model is Hempel's treatment of explanation in the non-probabilistic case.¹¹ He proposes that an explanation of a particular event consists, ideally, of a correct deductive-nomological (henceforth D–N) argument. There are law premises and particular-fact premises and no others. The conclusion

¹¹ For a full presentation of Hempel's views, see the title essay in his *Aspects of Scientific Explanation*.

says that the explanandum event took place. The argument is valid, in the sense that the premises could not all be true and the conclusion false. (We might instead define validity in syntactic terms. If so, we should be prepared to include mathematical, and perhaps definitional, truths among the premises.) No premise could be deleted without destroying the validity of the argument. The premises are all true.

Hempel also offers a treatment for the probabilistic case; but it differs significantly from his deductive-nomological model, and also it has two unwelcome consequences. (1) An improbable event cannot be explained at all. (2) One requirement for a correct explanation—"maximal specificity"—is relative to our state of knowledge; so that our ignorance can make correct an explanation that would be incorrect if we knew more. Surely what's true is rather that ignorance can make an explanation seem to be correct when really it is not. Therefore, instead of Hempel's treatment of the probabilistic case, I prefer to consider Railton's "deductive-nomological model of probabilistic explanation".¹² This closely parallels Hempel's D-N model for the non-probabilistic case, and it avoids both the difficulties just mentioned. Admittedly, Railton's treatment is available only if we are prepared to speak of chances—single-case objective probabilities. But that is no price at all if we have to pay it anyway. And we do, if we want to respect the apparent content of science. (Which is not the same as

¹² See Railton's paper of the same name. In what follows I shall simplify Railton's position in two respects. (1) I shall ignore his division of a D-N argument for a probabilistic conclusion into two parts, the first deriving a law of uniform chances from some broader theory and the second applying that law to the case at hand. (2) I shall pretend, until further notice, that Railton differs from Hempel only in his treatment of probabilistic explanation; in fact there are other important differences, to be noted shortly.

It is important to distinguish Railton's proposal from a different way of using single-case chances in a covering-law model of explanation, proposed in James H. Fetzer, "A Single Case Propensity Theory of Explanation," *Synthese* 28 (1974), pp. 171–98. For Fetzer, as for Railton, the covering laws are universal generalizations about single-case chances. But for Fetzer, as for Hempel, the explanatory argument, without any addendum, is the whole of the explanation; it is inductive, not deductive; and its conclusion says outright that the explanandum took place, not that it had a certain chance. This theory shares some of the merits of Railton's. However, it has one quite peculiar consequence. For Fetzer, as for Hempel, an explanation is an argument; however, a good explanation is not necessarily a good argument. Fetzer, like Railton, wants to have explanations even when the explanandum is extremely improbable. But in that case a good explanation is an extremely bad argument. It is an inductive argument whose premises not only fail to give us any good reason to believe the conclusion, but in fact give us very good reason to *disbelieve* the conclusion.

respecting the positivist philosophy popular among scientists.) Frequencies—finite or limiting, actual or counterfactual—are fine things in their own right. So are degrees of rational belief. But they just do not fit our ordinary conception of objective chance, as exemplified when we say that any radon-222 atom at any moment has a 50% chance of decaying within the next 3.825 days. If chances are good enough for theorists of radioactive decay, they are good enough for philosophers of science.

Railton proposes that an explanation of a particular chance event consists, ideally, of two parts. The first part is a D-N argument, satisfying the same constraints that we would impose in the nonprobabilistic case, to a conclusion that the explanandum event had a certain specified chance of taking place. The chance can be anything: very high, middling, or even very low. The D-N argument will have probabilistic laws among its premises—preferably, laws drawn from some powerful and general theory—and these laws will take the form of universal generalizations concerning single-case chances. The second part of the explanation is an addendum—not part of the argument—which says that the event did in fact take place. The explanation is correct if both parts are correct: if the premises of the D-N argument are all true, and the addendum also is true.

Suppose we have a D-N argument, either to the explanandum event itself or to the conclusion that it has a certain chance. And suppose that each of the particular-fact premises says, of a certain particular event, that it took place. Then those events are jointly sufficient, given the laws cited, for the event or for the chance. In a sense, they are a minimal jointly sufficient set; but a proper subset might suffice given a different selection of true law premises, and also it might be possible to carve off parts of the events and get a set of the remnants that is still sufficient under the original laws. To perform an act of explaining by producing such an argument and committing oneself to its correctness is, in effect, to make two claims: (1) that certain events are jointly sufficient, under the prevailing laws, for the explanandum event or for a certain chance of it; and (2) that only certain of the laws are needed to establish that sufficiency.

It would make for reconciliation between my account and the covering-law model if we had a covering-law model of causation to go with our covering-law model of explanation. Then we could rest assured that the jointly sufficient set presented in a D-N argument was a set of causes of the explanandum event. Unfortunately, that assurance is not to be had. Often, a member of the jointly sufficient set pre-

sented in a D–N argument will indeed be one of the causes of the explanandum event. But it may not be. The counterexamples are well known; I need only list them.

1. An irrelevant non-cause might belong to a non-minimal jointly sufficient set. Requiring minimality is not an adequate remedy; we can get an artificial minimality by gratuitously citing weak laws and leaving stronger relevant laws uncited. That is the lesson of Salmon's famous example of the man who escapes pregnancy because he takes birth control pills, where the only cited law says that nobody who takes the pills becomes pregnant, and hence the premise that the man takes pills cannot be left out without spoiling the validity of the argument.¹³

2. A member of a jointly sufficient set may be something other than an event. For instance, a particular-fact premise might say that something has a highly extrinsic or disjunctive property. I claim that such a premise cannot specify a genuine event; see "Events," in this volume.

3. An effect might belong to a set jointly sufficient for its cause, as when there are laws saying that a certain kind of effect can be produced in only one way. That set might be in some appropriate sense minimal, and might be a set of events. That would not suffice to make the effect be a cause of its cause.

4. Such an effect might also belong to a set jointly sufficient for another effect, perhaps a later effect, of the same cause. Suppose that, given the laws and circumstances, the appearance of a beer ad on my television could only have been caused by a broadcast which would also cause a beer ad to appear on your television. Then the first appearance may be a member of a jointly sufficient set for the second; still, these are not cause and effect. Rather they are two effects of a common cause.

5. A preempted potential cause might belong to a set jointly sufficient for the effect it would have caused, since there might be nothing that could have stopped it from causing that effect without itself causing the same effect.

In view of these examples, we must conclude that the jointly sufficient set presented in a D–N argument may or may not be a set of causes. We do not, at least not yet, have a D–N analysis of causation. All the same, a D–N argument may present causes. If it does, or rather

¹³ See Wesley C. Salmon et al., *Statistical Explanation and Statistical Relevance* (Pittsburgh: University of Pittsburgh Press, 1971), p. 34.

if it appears to the explainer and audience that it does, then on my view it ought to look explanatory. That is the typical case with sample D–N arguments produced by advocates of the covering-law model.

If the D–N argument does not appear to present causes, and it looks explanatory anyway, that is a problem for me. In Section III, I discussed three such problem cases; the alleged non-causal explanations there considered could readily have been cast as D–N arguments, and indeed I took them from Hempel's and Railton's writings on covering-law explanation. In some cases, I concluded that information was after all given about how the explanandum was caused, even if it happened in a more roundabout way than by straightforward presentation of causes. In other cases, I concluded that what was explained was not really a particular event. Either way, I'm in the clear.

If the D–N argument does not appear to present causes, and therefore fails to look explanatory, that is a problem for the covering-law theorist. He might just insist that it *ought* to look explanatory, and that our customary standards of explanation need reform. To the extent that he takes this high-handed line, I lose interest in trying to agree with as much of his theory as I can. But a more likely response is to impose constraints designed to disqualify the offending D–N arguments. Most simply, he might say that an explanation is a D–N argument of the sort that does present a set of causes, or that provides information in some more roundabout way about how the explanandum was caused. Or he might seek some other constraint to the same effect, thereby continuing the pursuit of a D–N analysis of causation itself. Railton is one covering-law theorist who acknowledges that not just any correct D–N argument (or probabilistic D–N argument with addendum) is explanatory; further constraints are needed to single out the ones that are. In sketching these further constraints, he does not avoid speaking in causal terms. (He has no reason to, since he is not attempting an analysis of causation itself.) For instance, he distinguishes D–N arguments that provide an "account of the mechanism" that leads up to the explanandum event; by which he means, I take it, that there ought to be some tracing of causal chains. He does not make this an inescapable requirement, however, because he thinks that not all covering-law explanation is causal.¹⁴

A D–N argument may explain by presenting causes, or otherwise giving information about the causal history of the explanandum; is it

¹⁴ See his *Explaining Explanation*, "A Deductive-Nomological Model of Probabilistic Explanation," and "Probability, Explanation, and Information."

also true that any causal history can be characterized completely by means of the information that can be built into D–N arguments? That would be so if every cause of an event belongs to some set of causes that are jointly sufficient for it, given the laws; or, in the probabilistic case, that are jointly sufficient under the laws for some definite chance of it. Is it so that causes fall into jointly sufficient sets of one or the other sort? That does not follow, so far as I can tell, from the counterfactual analysis of causation that I favor. It may nevertheless be true, at least in a world governed by a sufficiently powerful system of (strict or probabilistic) laws; and this may be such a world. If it is true, then the whole of a causal history could in principle be mapped by means of D–N arguments (with addenda in the probabilistic case) of the explanatory sort.

In short, if explanatory information is information about causal histories, as I say it is, then one way to provide it is by means of D–N arguments. Moreover, under the hypothesis just advanced, there is no explanatory information that could not in principle be provided in that way. To that extent the covering-law model is dead right.

But even when we acknowledge the need to distinguish explanatory D–N arguments from others, perhaps by means of explicitly causal constraints, there is something else wrong. It is this. The D–N argument—correct, explanatory, and fully explicit—is represented as the ideal serving of explanatory information. It is the right shape and the right size. It is enough, anything less is not enough, and anything more is more than enough.

Nobody thinks that real-life explainers commonly serve up full D–N arguments which they hope are correct. We very seldom do. And we seldom could—it's not just that we save our breath by leaving out the obvious parts. We don't know enough. Just try it. Choose some event you think you understand pretty well, and produce a fully explicit D–N argument, one that you can be moderately sure is correct and not just almost correct, that provides some non-trivial explanatory information about it. Consult any science book you like. Usually the most we can do, given our limited knowledge, is to make existential claims.¹⁵ We can venture to claim that there exists some (correct, etc.)

¹⁵ In *Foundations of Historical Knowledge*, Chapter III, Morton White suggests that "because"-statements should be seen as existential claims. You assert the existence of an explanatory argument which includes a given premise, even though you may be unable to produce the argument. This is certainly a step in the right direction. However it seems to underestimate the variety of existential statements that might be made, and also it incorporates a suspect D–N analysis of causation.

D–N argument for the explanandum that goes more or less like this, or that includes this among its premises, or that draws its premises from this scientific theory, or that derives its conclusion from its premise with the aid of this bit of mathematics, or I would commend these existential statements as explanatory, to the extent—and only to the extent—that they do a good job of giving information about the causal history of the explanandum. But if a proper explanation is a complete and correct D–N argument (perhaps plus addendum), then these existential statements are not yet proper explanations. Just in virtue of their form, they fail to meet the standard of how much information is enough.

Hempel writes "To the extent that a statement of individual causation leaves the relevant antecedent conditions, and thus also the requisite explanatory laws, indefinite it is like a note saying that there is a treasure hidden somewhere."¹⁶ The note will help you find the treasure provided you go on working, but so long as you have only the note you have no treasure at all; and if you find the treasure you will find it all at once. I say it is not like that. A shipwreck has spread the treasure over the bottom of the sea and you will never find it all. Every dubloon you find is one more dubloon in your pocket, and also it is a clue to where the next dubloons may be. You may or may not want to look for them, depending on how many you have so far and on how much you want to be how rich.

If you have anything less than a full D–N argument, there is more to be found out. Your explanatory information is only partial. Yes. *And so is any serving of explanatory information we will ever get*, even if it consists of ever so many perfect D–N arguments piled one upon the other. There is always more to know. A D–N argument presents only one small part—a cross section, so to speak—of the causal history. There are very many other causes of the explanandum that are left out. Those might be the ones we especially want to know about. We might want to know about causes earlier than those presented. Or we might want to know about causes intermediate between those presented and the explanandum. We might want to learn the mechanisms involved by tracing particular causal chains in some detail. (The premises of a D–N argument might tell us that the explanandum would come about through one or the other of two very different causal chains, but not tell us which one.) A D–N argument might give us far from enough explanatory information, considering what sort of information we

¹⁶ *Aspects of Scientific Explanation*, p. 349.

want and what we possess already. On the other hand, it might give us too much. Or it might be the wrong shape, and give us not enough and too much at the same time; for it might give us explanatory information of a sort we do not especially want. The cross-section it presents might tell us a lot about the side of the causal history we're content to take for granted, and nothing but stale news about the side we urgently want to know more about.

Is a (correct, etc.) D–N argument in *any* sense a complete serving of explanatory information? Yes in this sense, and this sense alone: it completes a jointly sufficient set of causes. (And other servings complete seventeen-membered sets, still others complete sets going back to the nineteenth century. . . .) The completeness of the jointly sufficient set has nothing to do with the sort of enoughness that we pursue. There is nothing ideal about it, in general. Other shapes and sizes of partial servings may be very much better—and perhaps also better within our reach.

It is not that I have some different idea about what is the unit of explanation. We should not demand a unit, and that demand has distorted the subject badly. It's not that explanations are things we may or may not have one of; rather, explanation is something we may have more or less of.

One bad effect of an unsuitable standard of enoughness is that it may foster disrespect for the explanatory knowledge of our forefathers. Suppose, as may be true, that seldom or never did they get the laws quite right. Then seldom or never did they possess complete and correct D–N arguments. Did they therefore lack explanatory knowledge? Did they have only some notes, and not yet any of the treasure? Surely not! And the reason, say I, is that whatever they may not have known about the laws, they knew a lot about how things were caused.

But once again, the covering-law model needn't have the drawback of which I have been complaining; and once again it is Railton who has proposed the remedy.¹⁷ His picture is similar to mine. Associated with each explanandum we have a vast and complicated structure; explanatory information is information about this structure; an act of explaining is an act of conveying some of this information; more or less information may be conveyed, and in general the act of explaining may be more or less satisfactory in whatever ways any act of conveying information about a large and complicated structure may be more or

less satisfactory. The only difference is that whereas for me the vast structure consists of events connected by causal dependence, for Railton it is an enormous "ideal text" consisting of D–N arguments—correct, satisfying whatever constraints need be imposed to make them explanatory, and with addenda as needed—strung together. They fit together like proofs in a mathematics text, with the conclusion of one feeding in as a premise to another, and in the end we reach arguments to the occurrence, or at least a chance, of the explanandum itself. It is unobjectionable to let the subject matter come in units of one argument each, so long as the activity of giving information about it needn't be broken artificially into corresponding units.

By now, little is left in dispute. Both sides agree that explaining is a matter of giving information, and no standard unit need be completed. The covering-law theorist has abandoned any commitment he may once have had to a D–N analysis of causation; he agrees that not just any correct D–N argument is explanatory; he goes some distance toward agreeing that the explanatory ones give information about how the explanandum is caused; and he does not claim that we normally, or even ideally, explain by producing arguments. For my part, I agree that one way to explain would be to produce explanatory D–N arguments; and further, that an explainer may have to argue for what he says in order to be believed. Explanation as argument versus explanation as information is a spurious contrast. More important, I would never deny the relevance of laws to causation, and therefore to explanation; for when we ask what would have happened in the absence of a supposed cause, a first thing to say is that the world would then have evolved lawfully. The covering-law theorist is committed, as I am not, to the thesis that all explanatory information can be incorporated into D–N arguments; however, I do not deny it, at least not for a world like ours with a powerful system of laws. I am committed, as he is not, to the thesis that all explaining of particular events gives some or other sort of information about how they are caused; but when we see how many varieties of causal information there are, and how indirect they can get, perhaps this disagreement too will seem much diminished.

One disagreement remains, central but elusive. It can be agreed that information about the prevailing laws is at least highly relevant to causal information, and *vice versa*; so that the pursuit of explanation and the investigation of laws are inseparable in practice. But still we can ask whether information about the covering laws is itself *part* of explanatory information. The covering law theorist says yes; I say no. But this looks like a question that would be impossible to settle, given

¹⁷ See *Explaining Explanation* and "Probability, Explanation, and Information."

Events*

240

Causation

that there is no practical prospect of seeking or gaining information about causes without information about laws, or information about laws without information about causes. We can ask whether the work of explaining would be done if we knew all the causes and none of the laws. We can ask; but there is little point trying to answer, since intuitive judgments about such preposterous situations needn't command respect.

I. INTRODUCTION

Events are not much of a topic in their own right. They earn their keep in the discussion of other topics: sometimes the semantics of nominalisations and adverbial modification, sometimes the analysis of causation and causal explanation. There is no guarantee that events made for semantics are the same as the events that are causes and effects. It seems unlikely, in some cases at least. A certain mathematical sequence converges. There is some entity or other that we may call the converging of the sequence. The sequence converges rapidly iff, in some sense, this entity is rapid. I have no objection to that; but I insist that the converging of the sequence, whatever it may be, is nothing like any event that causes or is caused. (The so-called "events" of probability theory are something else again—propositions, or properties of things at times.) My present interest is in events as causes and effects. Therefore I shall not follow the popular strategy of approaching events by way of nominalisations. Events made in the image of nominalisations are right for some purposes, but not for mine. When I introduce nominalisations to denote events, as I shall, it will not be analysis of natural language but mere stipulative definition.

* I am much indebted to discussions with Jonathan Bennett, Alison McIntyre, and Mark Johnston.