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

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.

Twenty-Three Events[*]

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.

In the two previous papers, I put forward several theses about causation and explanation. (1) Causal dependence is counterfactual dependence between distinct events. Event *e* depends causally on the distinct event *c* iff, if *c* had not occured, *e* would not have occurred—or at any rate, *e*'s chance of occurring would have been very much less than it actually was. (We must take care to use the right kind of counterfactuals: no backtrackers. See "Counterfactual Dependence and Time's Arrow," in this volume.) (2) Causation is the ancestral of causal dependence: event *c* causes event *e* iff either *e* depends on *c*, or *e* depends on an intermediate event *d* which in turn depends on *c*, or. . . . Causation without direct causal dependence is exceptional, but it occurs in cases of causal preemption. (See Postscript E to "Causation," in this volume.) (3) Any event has a causal history: a vast branching structure consisting of that event and all the events which cause it, together with all the relations of causal dependence among these events. (4) To explain why an event occurs is to give information about its causal history. Such information is inevitably partial. An explanation may specify part of the causal history of the explanandum event, or it may just provide structural information of one or another sort about the causal history. Goodness of explanation is governed by the pragmatic standards that apply to information-giving generally.

Since these four theses concern causation among events, their meaning cannot be entirely clear until I provide a theory of events to go with them. Not just any theory will do. If a theory posits too many distinct events, then many instances of counterfactual dependence between its allegedly distinct alleged events will clearly not be causal.[1] This difficulty will arise, for instance, on a theory that posits an abundance of distinct events to match the abundance of nonequivalent

predicates in nominalisations. A theory that allows unlimited Boolean combination of events also will generate alleged events that enter into relations of counterfactual dependence, but that do not seem intuitively suited to cause or to be caused. On the other hand, a sparse theory may posit too few events, forcing us to go beyond the events it countenances in order to complete our causal histories. This difficulty will arise, for instance, on a theory that limits itself to events falling under eventnouns of ordinary language: flashes, bangs, thumps, bumps, lectures, kisses, battles. . . . More generally, it will arise on a theory that provides no events to fill those stretches of time that we call "uneventful."

A theory that gives events unduly rich and fragile essences also will make trouble, as we shall see.

In this paper I shall consider what sort of theory of events I need to go with my theses about causation. If none could be found, that would be reason to reject what I say about causation. But I think a suitable theory can be found—or at least sketched—and I think it is a reasonably attractive theory in its own right. What other purposes it might serve, if any, I cannot say.

II. Events Are Properties of Spatiotemporal Regions

An event is a localised matter of contingent fact. It occurs. It is contingent that it occurs; no event occurs at every possible world. Hence we have nonvacuous counterfactuals about what would have been the case if a given event had not occurred, as we must if we are to place that event in a history of causal dependence. An event occurs in a particular spatiotemporal region. Its region might be small or large; there are collisions of point particles and there are condensations of galaxies, but even the latter occupy regions small by astronomical standards.

(Perhaps not just any region is a region in which an event can occur. A smallish, connected, convex region may seem a more likely candidate than a widely scattered part of spacetime. But I leave this question unsettled, for lack of clear test cases. If all of this year's VFL football comprises one big event, that event occurs in a scattered region; bits of it occur in various parks on various afternoons. But does it all comprise one event? Intuition is silent; and, so far as I can tell, the needs of my account of causation could be met either way.)

An event occurs in exactly one region of the world, if it occurs at all. If an event occurs in a region, it does not occur in any proper part of that region. The whole of the event occupies the whole of its region. Parts of it, but not the whole of it, may occur in parts of its region. Also, an event is unrepeated: it does not occur in two different regions of the world.

Thus an "annual event" such as the Grand Final is not an event in the sense of the present theory. As is only right and proper: the Grand Final does not cause or get caused *simpliciter*. It has different causes and effects in different years, which is to say that the different, unrepeated Grand Finals of the different years are what really have the causes and effects.

I distinguish occurring *in* a region and *within* a region. An event occurs within every region that includes the region in which it occurs; and it occurs in the region that is the intersection of all regions within which it occurs. We might also say that it *is occurring in* every region that is part of the region in which it occurs; and that it occurs in the region that is the mereological sum of all regions in which it is occurring.

To any event there corresponds a property of regions: the property that belongs to all and only those spatiotemporal regions, of this or any other possible world, in which that event occurs. Such a property belongs to exactly one region of any world where the event occurs, and there are some such worlds. It belongs to no region of any world where the event does not occur, and there are some of those worlds also. If a property of regions satisfies the conditions just stated, it may or may not correspond to an event. But at least it is *formally eligible* to do so.

By a *property* I mean simply a class—any class.[2] To have the property is to belong to the class. *All* the things that have the property, whether actual or merely possible, belong. My point in

using the word "property" is simply to emphasise that we are dealing with a class that may have otherworldly things, unactualised *possibilia*, among its members. (It might even, in the case of mathematical properties such as oddness, have *un*worldly members that are not part of any world.) The property that corresponds to an event, then, is the class of all regions—at most one per world— where that event occurs.

Two events can occur in exactly the same region. An electron's presence in a field can cause its acceleration; radiation of two frequencies can reverberate throughout the same cavity; two chemical reactions can go on in the same flask. More fancifully, there might be goblins made of a sort of matter that passes through our sort without any interaction, and a battle of goblins might occur in the very same region as this conference (the 1981 Conference of the Australasian Association of Philosophy, where this paper was first read).

But in each case it would have been possible for one of the two events to occur without the other. It would have been contrary to law, in some cases, but I take it that the laws of nature themselves are contingent. However, I can think of no plausible case of two events such that, necessarily, for any region, one occurs in that region iff the other does. Two such inseparables would be causally indistinguishable on a counterfactual analysis of causation, so it is hard to see how my treatment of causation could possibly need them both. I shall therefore take it that for any two events there is some region of some world where one occurs and the other does not. That region has the property corresponding to one event. It lacks the property corresponding to the other. So the two events correspond to two different properties. Our correspondence between events and properties of regions is therefore one-to-one.

A one-to-one correspondence is an opportunity for reduction, and I see no reason why events are needed as irreducible elements of being. Therefore, I propose to identify events with their corresponding properties. An event is a property, or in other words a class, of spatio-temporal regions. It satisfies my conditions of formal eligibility by containing one region each from some worlds, none from others, and never more than one from the same world. It occurs if and where and when there is a region that is a member of it.

Not just any property meeting the conditions given is to count as an event. I have said what kind of things events are—namely, formally eligible properties of regions—but not which things of that kind are events. The latter parts of the paper will address that question, though all I say will still fall sadly short of a precise necessary and sufficient condition for eventhood.

I am relying on the assumptions (1) that regions are individuals which are parts of possible worlds, and (2) that no region is part of two different worlds. Those assumptions are controversial, but I need not defend them here. Those who doubt them need only retreat to a more complicated set-theoretic construction of properties—as functions that

assign to each world the set of things that have that property at that world—whereupon what I say will go through essentially unchanged.

My proposal resembles that of Richard Montague, on which events are taken as certain properties of times.[3] The event occurs at a certain time at a certain world iff the property which is that event belongs, at that world, to that time. Thus he identifies the event with the property of being a time when that event occurs. I think my proposal has two minor advantages. (1) In view of Relativity, it is not altogether clear what sort of thing a time is. (2) Given that a Montague event occurs at a certain time at a certain world, we must work to recover the place where it occurs; given that one of my events occurs in a region of a world, its place is given immediately.

My proposal also resembles the suggestion that events might simply be identified with regions, or perhaps with regions including all that occupies them. That has been suggested by several authors, usually with some acknowledgement that what they suggest does not conform to normal usage.[4] We do usually think that two different events might occur in the very same region; not so, of course, if we identify events with their regions. If this conference is its region, and a battle of goblins is its region, and the conference and the battle are the same region, then the battle and the

conference are a single event. You might like the idea of treating events as regions, and yet you might insist on distinguishing these two events. If so, you might want to say that one event is the region *qua* conference venue, the other is the very same region *qua* battlefield. And you might want to say that they are identical, yet they are to be distinguished. In a similar frame of mind, you might want to say that Russell *qua* philosopher and Russell *qua* politician are identical, yet they are to be distinguished. You really should

not say such things: nothing may be distinguished in any way whatsoever from itself. Yet we may sympathise with your inclination, and provide for it legitimately as follows. Russell *qua* philosopher is Russell-taken-in-intension: the Russell of our world, taken together with the Russells of other worlds who are his philosophical counterparts. Likewise for Russell *qua* politician, except that instead of the philosophical counterparts we take the political counterparts. So far as this world is concerned, there is no difference between the two Russells-in-intension; their thisworldly members really are identical. But the otherworldly members differ, since our Russell has many philosophical-but-not-political and political-but-not-philosophical counterparts. So the two Russells-in-intension, taken entire as classes spread over many worlds, really do differ. Likewise for a pair of regions-in-intension such as this region *qua* conference venue and the same region *qua* battlefield: their thisworldly members are identical, but they differ by having different otherworldly members. My events are exactly such regions-in-intension, consisting of regions of many worlds united by suitable relations of similarity.

If events are properties, understood as classes with members from many worlds, then you might wish to say that an event exists whether or not it occurs: like a number, it exists necessarily, from the standpoint of every world, though it is not part of any world. Or you might instead say that an event exists at just those worlds where it occurs; or even that it does not fully exist at any world, since no world holds all of it. I think these are merely verbal questions. We may as well leave them unsettled; nothing hangs on them. Never mind whether it is contingent that an event *exists*; it is at any rate contingent that an event *occurs*. Also, I do not think it would be appropriate to deny the name "event" to those events that do not occur in our own world. After all, they are of a kind with the events that do actually occur. Others may prefer a different terminology, on which "events" are the ones that actually occur, and the rest are called something else. It matters not—the terminologies are intertranslatable.

III. Events Are Described Essentially and Accidentally

Events have their essences built in, in the form of necessary conditions for their occurrence. We may classify events by their essences, stating conditions that a region must satisfy if that event is to occur there. For

instance, an event is essentially a change iff, necessarily, that event occurs in a region only if something changes throughout that region.[5] Likewise an event is essentially a death iff, necessarily, that event occurs in a region only if someone dies throughout that region, and not throughout any larger one. Such a region will be the location of a final temporal segment of the victim, beginning just when he starts to die. (The vagueness of just when he starts to die infects our classification of events with vagueness, but there's no harm in that.) And so on for a wide range of essential classifications. These will include some made from single verbs, and others made from more complex predicate phrases—even infinitely complex ones. Thus we say what it would mean (whether or not it is ever true) that some event essentially is a vibrating-of-a-steel-gong-of-so-and-so-size-and-shape-at-so-and-so-frequency.

(When I use predicate phrases to define essential classifications of events, I am not making any claim of "conceptual priority," whatever that might mean. If the predicates in turn are definable in terms of the classification of events, we have nothing worse than a circle of interdefinables. Such

circles do not suffice to eliminate all the interdefined terms at once, of course, but they may nevertheless be useful collections of analytic truths.)

We can also say what it would mean (whether or not it is ever true) for an event to essentially involve Socrates. It does so iff, necessarily, it occurs in a region only if Socrates is present there—either the Socrates of our world, or else some otherworldly Socrates who is a counterpart of ours.[6] (The vagueness of the counterpart relation infects the classification with vagueness, but again we needn't mind.) Likewise, combining classification by predicates and by individuals, an event would be essentially a death of Socrates iff, necessarily, it occurs in a region only if Socrates dies throughout that region. Perhaps we should say (for reasons to be considered later) that the individuals essentially involved in events are not whole persisting people but temporal segments of them. But if so, that doesn't change the definition of involvement: for instance an event essentially involves a certain temporal segment of

Socrates iff, necessarily, that event occurs in a region only if that segment, or a counterpart of it, is present there.

We can also say what it would mean for an event to have its location, or a partial specification of its location, essentially. It essentially occurs in region R iff, necessarily, it occurs in a region only if that region is R, or a counterpart of R. It essentially occurs within the twentieth century iff, necessarily, it occurs only within the twentieth century, or a counterpart thereof.

Not only does an event have built-in necessary conditions for its occurrence; it has built-in necessary *and sufficient* conditions. That is just to say that there is a property that belongs to all *and only* the regions of this and other worlds where it occurs; and that is just to say that these regions comprise a class. If we could manage to express that property, and thus state necessary and sufficient conditions for the occurrence of an event, then not only could we classify that event by its essence; we could specify it uniquely. It would be the unique event such that, necessarily, it occurs in a region R iff . . . R. . . .

You might hope that an essential specification of an event could easily be extracted from the sort of nominalisation whereby we standardly denote it. Suppose we denote an event by a nominalisation: "the F-ing of A at T." Let f be the property expressed by the predicate F, let a be the individual denoted by A, and let t be the time denoted by T. (The denoting needn't be rigid.) The nominalisation denotes an event by way of the "constitutive" triple of f, a, and t; further, the occurrence of that event is somehow connected with the fact that property f belongs to individual a at time t.[7] (How does a property belong to an individual at a time? Perhaps because it is really a property of time-slices, or perhaps it is really a relation of individuals to times.) Then it is all too easy to assume that the triple gives us an essential specification of the denoted event. That is, we have the hypothesis that "the F-ing of A at T" denotes the event such that, necessarily, it occurs iff f belongs to a at t. (Iff so, then presumably it occurs in the region occupied by a at t.)

I think this will not do, at least not given the needs of a counterfactual analysis of causation. Sometimes, perhaps, an event can indeed be essentially specified in this way by means of a constitutive property,

individual, and time. But it is not so in general for the events we denote by nominalisations, and it is not so in general for causes and effects. The trouble is that an event with such a rich essence is a fragile thing. It is hard to change it without destroying it. It cannot occur at any but its constitutive time, it cannot involve any but its constitutive individual, and it cannot occur without something being an instance of its constitutive property. The causes and effects whereof we ordinarily speak are more robust than that.

The clearest difficulty concerns the alleged constitutive time. It is one thing to postpone an event, another to cancel it. A cause without which it would have occurred later, or sooner, is not a cause without which it would not have occurred at all. Who would dare be a doctor, if the hypothesis under consideration were right? You might manage to keep your patient alive until 4:12, when otherwise he would have died at 4:08. You would then have caused his death. For his death

was, in fact, his death at 4:12. If that time is essential, his death is an event that would not have occurred had he died at 4:08, as he would have done without your action. That will not do. (The point is due to Ken Kress. For further discussion, see Postscript E to "Causation," in this volume.)

Of course, we should not bounce off to the other extreme, and suppose that the death of the patient is an event such that, necessarily, it occurs iff he dies, never mind when and how. That would mean that the only way to cause someone's death would be to rob him of immortality, which is quite wrong also. Someone could die any of many different deaths, but not just any difference of time is enough to make the difference.

The alleged constitutive individual also is problematic. It is by no means clear that an event involving an individual always involves that individual essentially. Sometimes we are entitled to think of individuals as interchangeable parts. One member only of the firing squad got live ammunition, the rest fired blanks. The shooting was in fact done by Ted, but it could very well have been done by Ned instead. A cause without which someone else would have fired the fatal shot is not, or at any rate not clearly, a cause without which that very shooting—the one that was in fact a shooting by Ted—would not have taken place. It is not something on which that event depends.

Even the alleged constitutive property is not beyond suspicion. Perhaps any change, or any death, or any shooting, is such essentially. Perhaps not. But what if some much more specific, detailed predicate appears in the nominalisation? Sebastian strolled because he had plenty

end p.250

of time. Had he been delayed, the walking that was in fact a strolling might rather have been a striding. It might not even have been a walking, but rather a running. That is not to say, not clearly, that it would not have occurred at all.

So, while it is clear enough what it would mean to specify events essentially, often that does not seem to be what we really do. At any rate, it is not what we do when we specify events by means of our standard nominalisations.

Indeed, it may be no easy thing to refer to events by means of essential specifications of them. It would be downright impossible, if the event occurs in one but not both of two absolutely indiscernible regions of some world; and any event that could occur in a world of eternal recurrence is an event that occurs in two such regions. We might restrict our ambitions, ignore such especially troublesome worlds, and hope we could state a condition for the occurrence of an event which would be necessary and sufficient so far as the better-behaved worlds are concerned. That would not quite be an essential specification, but it would approximate to one. Even that much would, I fear, be a tall order; though the more we restrict ourselves, the easier it gets. Ignoring the worlds infected by indiscernibility might be only a beginning. If we restrict our attention to a small range of worlds sufficiently similar to actuality, we might have some hope of success at stating a condition that would be necessary and sufficient, so far as worlds in the range were concerned, for occurrence of a certain event. Of course, it would not distinguish events that differ only with respect to regions of worlds outside the range. Thus it could not be used to specify one event determinately. But often we find it tolerable to leave some indeterminacy in our specifications of things. Where exactly does the outback begin? Nobody knows; not because it's a secret, but because we've never bothered to settle exactly what "the outback" denotes. And yet we know, near enough, what we're talking about. It might be the same way with our best feasible approximations to essential specifications of events: the specification might be ambiguous between many events that more-or-less coincide throughout nearby worlds but differ at more distant worlds. That might be near enough to determinate reference to meet (some of) our needs. We might specify events that way; but then again we might not. Often, for instance when we denote events by means of the standard nominalisations, our specifications do not even approximate to being essential.

Likewise, while it is clear enough what it would mean to classify events essentially, it seems that very often we do not do that either. We

end p.251

specify and we classify actually occurring events in part by descriptions that fit these events accidentally. The event does fit the description, but that very event might have occurred without

fitting the description. The event is a class consisting of one region of this world together with various regions of other worlds wherein the event might have occurred. What goes on in the former region fits the description, but what goes on in some of the latter regions does not. The description may imply something about the essential classification of the event it describes, but it is not exhausted by that information.

There are many ways an event might be accidentally classified, and I cannot hope to give a complete inventory. (1) For events, as for anything else, we can always hoke up thoroughly artificial descriptions: "the event that is the Big Bang if Essendon will win the Grand Final, the birth of Calvin Coolidge if not." (2) We might classify it in part by its causes or its effects: "Fred's sunstroke," "what Fred did to bring it about that the window is open." (3) We might classify it with reference to its place in a conventional system: "Fred's signalling for a left turn," "Fred's signing of the cheque." (4) We might conjoin an accidental circumstance to an essential classification. A certain famous event was essentially a fiddling, let us assume, but only accidentally was it a fiddling while Rome burned.

I have already suggested that classifications by "constitutive" triples may be accidental, indeed that all three terms of the triple may be inessential. (5) We can classify or specify an event by its time, or more generally by the (exact or approximate) location of the region in which it actually occurs, even if it could have occurred at a somewhat different time and place. (6) Though there may be some events that involve "constitutive individuals" essentially, I have argued that others—as in the case of the firing squad—involve individuals accidentally. (7) Likewise for "constitutive properties." I don't deny that some event with a richly detailed essence might be essentially a strolling, so that necessarily it occurs only in a region wherein someone strolls. But some less fragile event might be only accidentally a strolling; it might be a strolling that could have been a striding. At the end of the next section, I shall consider the relation between essential and accidental strollings.

(8) If an event essentially involves one individual, it may thereby accidentally involve another. Suppose an event essentially involves a certain soldier, who happens to belong to a certain army. This event cannot occur in regions where there is no counterpart of that soldier, but it can occur where there is a counterpart of the soldier who does

not belong to a counterpart of that army. Then the event accidentally involves that army, by way of its soldier. Similarly, suppose an event essentially involves a certain temporal person-segment, which is accidentally a segment of Socrates; that would be one way for an event to accidentally involve Socrates.

(9) Another possibility for accidental classification turns on nonrigid designation of properties. I persist in thinking that "heat" nonrigidly designates whatever phenomenon it is that occupies a certain role and presents itself to us by causing certain manifestations.[8] In fact, this is molecular motion, but it might have been something else. A world where caloric fluid causes those manifestations is a world where the hot things are the ones with lots of caloric fluid. Then "the loss of heat by the poker" may denote an actually occurring event that is essentially a decreasing of molecular motion, and is only accidentally a loss of heat. This same event might have occurred at a world where caloric fluid is what presents itself as heat, in a region where the poker absorbed caloric fluid while its molecular motion decreased.

In any case of accidental classification or specification, the event actually described is one that might have occurred without fitting the given description. Whenever that is so, we must take care with our causal counterfactuals. Consider whether the event of a certain description would or would not have occurred under some counterfactual supposition. It is one thing to say that the event itself would not have occurred; it is a different thing to say that no event fitting the given description would have occurred. For the event might have occurred without fitting the description; or not that event, but another event fitting the description, might have occurred.

Many authors, most prominently Davidson, have noted that sentences which do not explicitly mention events often are equivalent to sentences which assert that there occurs (or exists) an event of such-and-such description.[2] We must at least agree that many such equivalences

hold, whether or not we regard them as somehow revealing underlying logical forms, and whether or not we think they are delivered wholesale in accordance with some nice general schema. The present treatment of events should not be expected to deliver them wholesale, as witness the mathematical sequence that converges although there occurs no event—in the sense I have in mind—which is its converging. But many such equivalences do hold, and these should guide us in considering the variety of ways in which events are classified. One reason to insist on accidental classification of events is that with essential classification alone, we do not get the proper equivalences. We get one direction easily enough; if there occurs an event that is essentially an *F*-ing, then in some region of this world an event occurs such that, necessarily, it occurs in a region only if something *F*-s there; so something *F*-s. But the converse direction will often fail. If something *F*-s, it does not in general follow that there occurs an event that is essentially an *F*-ing. The most that follows (and even that not in full generality, as witness the converging sequence) is that there occurs an event which, in some way or other, essentially or accidentally, is an *F*-ing. If Nero fiddles while Rome burns, I agree that an event occurs which is essentially a fiddling; and an event occurs which is accidentally a fiddling while Rome burns, since the aforementioned fiddling does occur while Rome burns. But for reasons to be considered later, I doubt that there occurs any event which is *essentially* a fiddling while Rome burns, and which could not have occurred under happier circumstances.

The foregoing discussion contributes only in a negative way to answering the question which ones among the formally eligible properties of regions are the events. We cannot answer that question correctly by first investigating our event descriptions, then taking events to be those properties of regions that such descriptions can specify essentially; because our specifications of events may be largely accidental.

IV. Events Stand in Logical Relations

Let us say that event *e implies* event *f* iff, necessarily, if *e* occurs in a region then also *f* occurs in that region. Considered as classes, event *e* is a subclass included in class *f*.

Just because we can define a relation, it doesn't follow that there are any instances of it. A theory of events might include the thesis that no events imply other events. That was my own view until recently, but now I think it should be rejected.

John says "Hello." He says it rather too loudly. Arguably there is one event that occurs which is essentially a saying-"Hello" and only accidentally loud; it would have occurred even if John had spoken softly. Arguably there is a second event that implies, but is not implied by, the first. This event is essentially a saying-"Hello"-loudly, and it would not have occurred if John had said "Hello" but said it softly. Both events actually occur, but the second could not have occurred without the first.

We have two descriptions: "John's saying 'Hello' " and "John's saying 'Hello' loudly." But it does not follow from this alone that we have two events to describe. The second description as well as the first might denote the first event, since the second description might describe the first event in part accidentally. Alternatively, the first description as well as the second might denote the second event, since the first description might describe the second event by less than the whole of its essence. Indeed, even if there are two different events, it still does not follow that one description denotes one and the other denotes the other. If both descriptions are somewhat vague or ambiguous, it could be that both denote both.

The real reason why we need both events, regardless of which description denotes which, is that they differ causally. An adequate causal account of what happens cannot limit itself to either one of the two. The first event (the weak one) caused Fred to greet John in return. The second one

(the strong one) didn't. If the second one had not occurred—if John hadn't said "Hello" so loudly—the first one still might have, in which case Fred still would have returned John's greeting. Also there is a difference on the side of causes: the second event was, and the first wasn't, caused *inter alia* by John's state of tension.[10]

We have two different events, causally distinguished. *Different*, but we needn't count them as *distinct*. "Distinct" does not mean "non-identical." I and my nose are not identical, but neither are we distinct. My nose is part-identical to me, identical to part of me. There is a clear sense in which our second event is part of the first: the subclass is part of the class, they are neither identical nor distinct. (Confusingly, by the inverse variation of intension and extension, there is also a sense in which the first event is part of the second.) Indeed, we dare not count the two as distinct. For their distinctness, plus my theses about causation, would together imply what is surely false: that the first event causes the second. For if the first had not occurred, then the second which implies the first would not have occurred either. Here is a case of non-causal counterfactual dependence—but *not* between distinct events. We may take it as a general principle that when one event implies another, then they are not distinct and their counterfactual dependence is not causal.

There is a persuasive intuition—I was long persuaded by it—that it is wrong to count both the first and the second event because if we do, we count something twice over. I now think that we do this intuition sufficient justice when we say that the first and second events, though not identical, also are not distinct. Compare the equally persuasive intuition that it is double-counting to include both atoms and molecules in our inventory of being—an adequate answer is that the molecules and their atoms are not distinct.

There might be two occurrent events that are both implied by some third occurrent event, but that are independent of each other. Like this: the first event is essentially a saying-"Hello"-loudly and is accidentally abrupt; the second is essentially a saying-"Hello"-abruptly, and is accidentally loud; the third event, which implies both the first and the second, is essentially a saying-"Hello"-loudly-and-abruptly. In this case also, the first and second should not be counted as distinct. Therefore they are not eligible to stand in causal dependence.

But we must beware. Suppose we had an occurrent event with a very rich essence: the unit class of a region of our world. It would imply all the events that occur in that region, none of them would then be distinct, and they would be ineligible to depend causally on one another. That is wrong. If this conference and the battle of goblins occurred in the same region, they would nevertheless be intuitively distinct, and they might perhaps stand in causal dependence. The acceleration of the electron does depend causally on its presence in the field. Therefore there must be a limit on how rich the essence of a genuine event can be.

Mere unit classes of regions are ruled out, and I should think a good deal else besides.

I have argued that we might have two events, one implying the other, such that one is essentially a saying-"Hello"-loudly, and the other is only accidentally a saying-"Hello"-loudly. I suggest that this illustrates one important way in which events come by their accidental classifications. We have, so to speak, a more and a less detailed version of what happens in a region. Both are occurrent events. The more detailed version has a richer essence; the otherworldly regions included in it are fewer and less varied; it is more tightly unified by similarity; there is less variety in the ways it could have occurred. The more detailed version is one, but only one, of the ways in which the less detailed version could have occurred. But it, unlike alternative more detailed versions, happens to be the way the less detailed version actually did occur. In such a case, the essential classification of the occurrent more detailed version carries over to become an accidental classification of the less detailed version. Likewise, I said above that some events are essentially strollings and some are accidentally strollings. The event that is accidentally a strolling is so because it is implied by another event, its more detailed version, which actually occurs and which is essentially a strolling. The accidental strolling has alternative more detailed versions; for instance, it is also implied by a non-occurrent event which is essentially a striding. If that one of the implying

events had been the one that occurred, then the implied event would have been accidentally a striding rather than a strolling.

(Above, I noted this implication: if there occurs an event that is essentially an *F*-ing, then it follows that something *F*-s. Now we can drop the "essentially," though subject to a proviso. For if there occurs an event that is accidentally an *F*-ing, provided it comes by this accidental classification in the way just discussed, then there occurs another event that is essentially an *F*-ing, and again it follows that something *F*-s.)

Similarly, if some occurrent event essentially involves some individual, then an event that it implies—its less detailed version—accidentally involves that individual. Consider the shooting that was done by Ted, but might have been done by Ned. How does it involve Ted? Perhaps (but this is not the only possible answer) because it is implied by an occurrent event, its more detailed version, that was a shooting essentially done by Ted.

Can it be, as I thought for a time, that accidental classifications of events *always* work this way? Can we say in general that they are

<div align="right">

end p.257
</div>

essential classifications of the occurrent more detailed version, carried over via implication? If so, then every accidental classification of one event would have to be an essential classification of another. But some classifications seem so very accidental that no event could have them essentially. Consider accidental classifications in terms of circumstances. There is an event that is accidentally a fiddling while Rome burns, but I doubt that any event is essentially a fiddling while Rome burns. And the example can be made even more extreme. There is an event that is accidentally classifiable as follows: it is a fiddling in the presence of a boy whose grandson will first set foot on the moon. Surely no event is essentially *that*! So we get no unified theory; there must be other ways for events to come by their accidental classifications. For accidental classifications in terms of circumstances, at any rate, it is no mystery how they manage to do it.

Again we would be in trouble if we had events with overly rich essences. If the unit classes of regions were counted as events, to take the worst case, then the accidental descriptions of events that happen to occur in the same region would coalesce. If this conference occurs in the same region as a battle of goblins, it would follow that the conference *is*, albeit accidentally, a battle of goblins. For it would be implied by an occurrent event—namely, the unit class—that is so essentially. That would never do.

V. Events Have a Spatiotemporal Mereology

We have seen how events may be, in a sense, logical parts of one another. If events are classes, as I propose, then they have a mereology in the way that all classes do: the parts of a class are its subclasses.

However, there is a second sense in which events have a mereology, and that will be our business in this section. Regions may be spatio-temporal parts of one another; events are classes of regions; the mereology of the members carries over to the classes, giving us a sense in which events also may be spatiotemporal parts of one another. Each of Sebastian's steps is a spatiotemporal part of his stroll; so is the entire half-stroll performed by the left half of him. Small events that occur in subregions are parts of the big event that occurs in the big region.

Let us say that event *e* is *essentially part* of event *f*, iff, necessarily, if *f* occurs in a region, then also *e* occurs in a subregion included in that region. (Not necessarily a proper subregion. Therefore an implied event is, in the sense of this definition, essentially part of the implying

<div align="right">

end p.258
</div>

event; though in another sense, the subclass sense, the implying event is part of the implied one. We have already noted this ambiguity.) However, events need not have their spatiotemporal parts essentially. Sebastian's stroll might have consisted of more steps, or fewer, or the same number but not the very same steps. A war might have consisted of different battles, though the

scope for difference is limited. We need to provide for accidental, as well as essential, spatiotemporal mereology of events.[11]

To do so, we may imitate the proposal of the previous section: let the essential mereology of the more detailed occurrent versions carry over via implication to become the accidental mereology of the less detailed versions. Let us say that occurrent event *e* is *part* of occurrent event *f* iff some occurrent event that implies *e* is essentially part of some occurrent event that implies *f*. This covers not only the case in which *e* is essentially part of *f*, but also the case in which *e* is only accidentally part of *f*, and *f* could have occurred without having *e* as a part.

When one event is part of another, whether essentially or accidentally, they are not identical but they are not distinct either. (Again, they are partly identical.) The same is true when two events share a part in common though neither is a part of the other. Kim gives the case of someone who writes "Larry," and as part of that event writes "rr."[12] I add that he writes "Larr" and he writes "rry," these being two overlapping events. No two of these four events are distinct. As in the case considered previously of events related by implication (which case is subsumed under the present one), events that are not distinct cannot stand in causal dependence. If the writing of "rr" had not occurred, the writing of "Larry" would not have, but that does not make those two events be cause and effect. Nor can the whole of the writing of "Larr" be a cause on which the writing of "rry" depends, though the first part of it well might be.

(Often we do say loosely that event *c* causes event *e* when what's true is that one or more parts of *c* cause one or more parts of *e*. Thus we might speak of some prolonged self-perpetuating process as its own cause, when really there is no self-causation and it is earlier parts that cause later parts. See Postscript A to "Causation," in this volume.)

Once more we would be in trouble if we had events with overly rich

essences. If the unit classes of regions were counted as events, it would turn out that whenever one event occurs within the region in which another occurs, the former event is part of the latter.

Once we have defined the part-whole relation for events, we can go on as usual to define other mereological notions. Thus events overlap iff they have some event as a common part; an atomic event is one that has no events except itself as parts; an event *e* is the mereological sum of events $f_1, f_2, \ldots$ iff *e* overlaps all and only those events that overlap at least one of the *f*'s; and so on. Then we can ask what principles this mereology of events obeys. Is it like the unrestricted mereology of individuals, in which several individuals always have another individual as their sum? Or is it like the restricted mereology of chairs, in which several chairs seldom, if ever, have another chair as their sum? Or is it in between? I suggest that events are at any rate more amenable to summation than chairs are: a war may be the sum of its battles, a conference may be the sum of its sessions. But I leave open the question whether several events, however miscellaneous, always have another event as their sum. If there is unrestricted summation, then there can be no limit on how large and disconnected and disunified an event may be; whereas if events must have some unity to them, then some attempted summations would fail to yield a genuine event. (Maybe they yield a property of regions which is formally eligible, but not an event.) It is hard to find arguments to settle the question. Our events are meant to serve as causes and effects; but it seems hard to tell when we can be content to say only that several events are joint causes and separate effects, and when we must also insist on a single event that is their sum.

Another question is this: given any subregion of the region in which an event occurs, is there a part of that event which occurs in that subregion? I am inclined to think that there is in the case of a suitable subregion—one with boundaries that are reasonably simple in shape, or that match the boundaries of something within the region—but not in the case of an arbitrary subregion. But again I cannot find arguments to settle the question.

VI. The History of Events Is the Whole of History

Suppose we are given the complete history of the world's events: exactly which events occur in exactly which regions, throughout all of space and time. Then no historical information is lacking. No two

possible worlds could be exactly alike in their histories of events, yet unlike elsewhere in their histories of matters of manifest particular fact. The total history of events implies every historical truth and contradicts every historical falsehood.

It must be so, given the four theses about causation and explanation that I took as my starting point. To explain is to give information about the causal history of the explanandum event, and that history is a structure of causally related events. If history is a patchwork of events and nonevents, and if the nonevents are not implicitly given by the history of events, then the nonevents are left out of causal histories. Then they never enter into the explanations of events. But is there really anything in history that is thus isolated, that never plays any part in explaining why any event took place? I do not see what it could be.

Not all events involve change. We cannot afford to count the unchanges as nonevents, for the unchanges may be needed to complete causal histories. Indeed, the causal history of an event—an uncontroversial event, an abrupt change—might consist entirely of unchanges. An isolated particle has existed from all eternity; it is unstable, and has at every moment a chance of decaying; eventually it does decay. The decay was caused (probabilistically) by the previous presence, at all earlier times, of the unchanging particle. The causal history that explains this event is entirely changeless. It is a thoroughly uneventful course of events.

The need for unchanges as events is urgent in connection with causal theories of perception, memory, persistence over time, and so on. The causal chains required by such theories often will consist simply of something continuing to exist: a travelling signal, a memory trace, a surviving person or a persisting lump of matter. The need is urgent also in connection with causal preemption. The preempting cause may cause its effect not because the effect depends directly on the cause, but because the effect depends causally on some intermediate event which in turn depends causally on the preempted cause. (See Postscript E to "Causation," in this volume.) If we could not count unchanges, there would be realistic cases in which we could find no suitable intermediate event.

Terminology is not the issue. If it is abuse of language to call unchanges "events," so be it. The point is that we must have them as causes and effects.

This section does not contribute directly to answering the question which among the formally eligible properties of regions are events. But it is indirectly relevant: no adequate demarcation of the events can be

too restrictive, else it leaves us with not enough events to make up the whole of history.

Given the history of events, the whole of history is implicitly given; but not, perhaps, the whole truth about contingent matters. It is an independent, and difficult, question whether two worlds exactly alike in their histories would have to be exactly alike in every way: in their chances, their laws, their modal truths and counterfactuals, their causal relations, and so on. (See the introduction to this volume.) My present point does not touch that question. I say just that total history supervenes on the history of events, whatever else may or may not supervene in turn on total history.

VII. Events Are Predominantly Intrinsic

Consider the alleged event such that, necessarily, it occurs in a region iff (1) that region is located at a certain time *t*, (2) Xanthippe is present in that region, and (3) at time *t*, someone dies who has been married to her until that time.[13] I shall call this alleged event *the widowing of Xanthippe*, since I need a name for it; but I do not mean to suggest (and do not believe) that this is what the phrase would ordinarily denote. (Maybe it would denote the death of Socrates; or maybe no event at all, but rather a certain fact, that is a certain true proposition.) Condition (3) is entirely

extrinsic. It has nothing to do with the qualitative character of the region in question, and everything to do with what goes on at other times and places. The other conditions are at least partly intrinsic. But they do not go very far toward picking out those regions where the alleged event allegedly occurs.

We can devise a more extreme, but also more artificial, case in which a supposed event is almost purely extrinsic in its essence. Begin with some genuine event: this conference, let us say. Define its *centennial* to be that event—if such it be—such that, necessarily, it occurs in a region R iff the original event occurs in the region which results from shifting R exactly 100 years backward, holding fixed its place, size, and shape. A centennial in this artificial sense is not at all intrinsic. Whether it occurs in a region is completely independent of what things are there, how they are arranged, what they are like, and what they are doing. It is quite unlike a centennial in the ordinary sense: a genuine event wherein some previous event is remembered and celebrated. The qualitative

end p.262

character of the region could be anything. (With one exception: it must have a size and shape which the original event could possibly have.)

We can define an *intrinsic* property of a region as one such that, whenever two possible regions are perfect duplicates, the property belongs to both or neither. Likewise a *purely extrinsic* property is one such that, for any possible region, there is some possible region which is a perfect duplicate of it and has the property.[14]

I think there are no such events as the widowing of Xanthippe or the centennial of this conference. We have the properties of regions, right enough, and they are formally eligible. But they are not events because they are purely or predominantly extrinsic, whereas the properties of regions that are genuine events are predominantly intrinsic.

Extrinsic events—or, more generally, events not predominantly intrinsic—are objectionable on three counts. (1) They offend our sense of economy. We would seem to count the death of Socrates twice over in our inventory of events, once as itself and again as (what I am calling) the widowing of Xanthippe. Still more clearly, the centennial of this conference is but a shadow of the conference itself. (2) They stand in relations of noncausal counterfactual dependence to those genuine events in virtue of which they occur. Without the death of Socrates, the widowing of Xanthippe would not have occurred. (She might still have been widowed sooner or later. But recall that the widowing of Xanthippe, as I defined it, had its time essentially.) Without this conference, its centennial would not occur. What's worse: without its centennial, the conference would not occur. None of these is a genuine case of causal dependence. (3) They also stand in relations of noncausal counterfactual dependence to other genuine events, events logically independent of them. Without the widowing of Xanthippe, the subsequent cooling of Socrates's body would not have occurred. (For in that case he would not have died when he did.) Without the centennial of this conference 100 years hence, our homeward departures a few days from now would not occur. These also are not genuine cases of causal dependence. Instantaneous and backward causation are not so very easy!

The first two objections might be answered by claiming that the events in question are in some sense not distinct. It is a bit hard to say

end p.263

why not: in the case of the death of Socrates and the widowing of Xanthippe, either could have occurred without the other. (They are logically, though not counterfactually, independent.) Further, they occur in nonoverlapping regions. But suppose that some suitable sense of distinctness could be found. That would do nothing at all to answer the third objection, which is therefore the decisive one.

Should I say that genuine events must be entirely, not just predominantly, intrinsic? I think not. For if an event were an entirely intrinsic property of regions, then it would have to occur in any duplicate of any region in which it occurs. Suppose it is an event that could occur within some epoch of some world of eternal recurrence—that's not a very strong condition, especially not on the hypothesis that events are entirely intrinsic. But every region within an epoch has duplicates within

all the other epochs. So this event occurs many times over in a single world, *contra* our stipulation that we are talking about particular events. That stipulation applies even to a world of eternal recurrence; because surely we want to distinguish the events of different epochs, indiscernible though they may be. A similar problem arises from duplication of regions within more ordinary worlds. It may be hard for big regions to be exactly alike in what goes on in them; but remember that submicroscopic events go on in submicroscopic regions, and electron-sized regions will have less opportunity to differ than larger regions do. We don't want a collision of two electrons to be one event that occurs in all regions of the same world where two electrons collide in just that way; rather, we want to distinguish different collisions that occur once each. If the regions do not differ intrinsically, then an event that occurs in one but not all of them cannot be entirely intrinsic. It can be predominantly intrinsic, and that is all I should require.

The rejection of overly extrinsic events allows me to return to some unfinished business. Earlier, I asked how an event comes to be accidentally classified as a fiddling while Rome burns; not, I said, because it has an occurrent more detailed version that is essentially a fiddling while Rome burns. Now I can give the reason: the alleged event that is essentially a fiddling while Rome burns would be too extrinsic.

I also avoided committing myself to the existence of genuine events which essentially involve Socrates; that is, which cannot occur except in a region where our Socrates, or a counterpart of him, is present. Therefore, I also could not commit myself to the existence of events that involve Socrates accidentally by having occurrent more detailed versions that involve him essentially. Now I can give the reason: being a counterpart of Socrates is rather an extrinsic matter. Counterparts are

united by similarity, and the similarity in question may be largely extrinsic. This is so especially on counterpart relations which stress match of origins: Socrates might have been—some counterpart of him is—a dimwitted and taciturn politician, most noted for his good looks and slender build; but only if his origins had been more or less exactly as they actually were. I am less of an enthusiast for match of origins than most philosophers of modality; nevertheless, I must agree that it often carries a lot of weight. And if similarities in respect of one's role in society and in history also carry weight in making counterparts, as I think they sometimes do, those too are extrinsic. The origins and role of Socrates are not intrinsic even to Socrates taken as a whole. Still less are they intrinsic to the temporal part of him whereby he is present in a region and involved there in some event. So the property of having present in it a counterpart of Socrates is not an intrinsic property of a region; so an alleged event that essentially involved Socrates would be to that extent extrinsic. I am not sure whether it would be extrinsic enough to deserve rejection.

(One worry about it is needless, I think. We might fear that any extrinsic element in an event will give rise to cases of noncausal counterfactual dependence. Suppose that Socrates is essentially the outcome of a certain event of conception, and that his death is essentially his. Then his death is essentially preceded by his conception. Had his conception not occurred, his death would not have—because then no death could have been *his*. But this time the dependence does seem causal, so we have no problem. Is there perhaps a new problem: causal dependence between *in*distinct events? No; there is a sense in which the death implies the conception, and I did say that when one event implies another in a certain sense they are not distinct and not eligible to stand in causal dependence, but the two senses are not the same. Nothing I said detracts from the distinctness, and eligibility for dependence, of two events that occur in nonoverlapping regions.)

Can we afford to reject events that essentially involve Socrates? That depends on whether we can do without them. Certainly some events involve Socrates at least accidentally. But can we find any way for an event to involve Socrates accidentally if none involves him essentially? Of course it will not do to say that he is involved just by being present in the region where the event actually occurs—what if those goblins are present in the region of this conference? Hence my rather cumbersome suggestion about involvement via a temporal segment. It is not necessary that counterpart relations must work in the same way for the parts as for the whole, or that the counterparts of the parts must be

parts of the counterpart of the whole. I find it somewhat plausible that a counterpart relation for temporal segments of people could work more by intrinsic similarity and less by origins or roles than a counterpart relation for whole people does. If so, I have less difficulty with an event that essentially involves a person-segment than with one that essentially involves a person. Now suppose we have a death that involves a certain person-segment. (Whether essentially or accidentally doesn't matter; if we have one that involves it accidentally, we have another that involves it essentially.) Now suppose that in fact this segment is part of Socrates. Accidentally so: not all its counterparts are parts of his. Now at last we've got Socrates involved in his own death, and in a way that bypasses any unduly extrinsic events. I wouldn't say no to a simpler method!

VIII. Events Are not Disjunctive

Let us call event *e* the *disjunction* of events $f_1, f_2, \ldots$ iff, necessarily, event *e* occurs in a region iff either f_1 or f_2 or $\ldots$ occurs there. (There may be finitely or infinitely many disjuncts.) Equivalently, *e* is the disjunction of the *f*'s iff each of the *f*'s implies *e*, and *e* implies any other event of which the same is true. Considered as classes, a disjunction of events is simply their union.

I do not deny that some events are disjunctions of others. An event that is essentially a stamping might, for instance, be the disjunction of one event that is essentially a stamping-the-left-foot and another that is essentially a stamping-the-right-foot. But I do not think that just any class will have a disjunction. The disjuncts must not be too miscellaneously varied. In calling an alleged event *disjunctive*, as opposed to saying that it is the disjunction *of* such-and-such, I mean that it would be the disjunction of some disjuncts that are overly varied. An example might be the supposed disjunction of one event that is essentially a walking and another that is essentially a talking.

Disjunctive events are *prima facie* open to the same three objections that I raised against extrinsic events. (1) They offend against economy. To count both disjunction and disjunct looks like counting the same thing twice. But to this an adequate reply is that the two are not distinct, by definitions already given. (2) They stand in relations of noncausal counterfactual dependence with their logical relatives, namely their disjuncts. Without the disjunction, no disjunct could have occurred. Without the occurrent disjunct, the disjunction would not

have occurred (except in special cases where some other disjunct would have occurred instead). It is an adequate reply that these are not cases of counterfactual dependence between distinct events. (It is just as well that these two objections fail. They threaten to prove too much— namely, that no event is *ever* the disjunction of others.) (3) They stand in relations of noncausal counterfactual dependence with other events that are not their logical relatives, and are clearly distinct from them. This objection is the decisive one.

Fred talks, and his talking causes Ted to laugh. Suppose that besides Fred's talking there is another event, the disjunctive event of Fred's talking-or-walking. Without it, Fred's talking would not have occurred, and neither would Ted's laughing. So this disjunctive event also causes Ted to laugh. That is intuitively wrong. No such event causes Ted's laughing, or anything else. Given the theses I took as my starting point, that can only be because there is no such event. Hence disjunctive events are to be rejected.

This may seem too hasty. It may seem that there are some disjunctive events that we dare not lose, because they are indispensable as causes and effects. Suppose that a certain poker cools down; that is, it loses heat. Suppose also that I am right to say that "heat" is not a rigid designator; heat is whatever property it is that occupies a certain causal role, and so might be one property or might be another. So there are many quite different ways that the poker might lose heat, depending on what sort of world it is in. (It or its counterpart.) Its molecular motion might decrease, in a world where molecular motion is what occupies the role; or it might lose caloric fluid, in a world where caloric fluid is what occupies the role; or. . . . So it seems that losing heat is quite a disjunctive

affair; and what's worse, extrinsic, since whether one property or another occupies the heat-role depends on what goes on throughout the world in question, not just on the region of it where the poker is. All the same, isn't the loss of heat by the poker a perfectly good event? Isn't it a cause of the poker's slight contraction, and an effect of my taking the poker out of the fire?

I agree that the loss of heat by the poker is a perfectly good event. But it is not disjunctive, so we needn't fear to lose it if we reject disjunctive events. For it is not the disjunction of one event which is essentially a decreasing of molecular motion in a world where that occupies the role, and another which is essentially a loss of caloric fluid in a world where that occupies the role, and. . . . That would be an event that was *essentially* a losing of heat. (And that had no relevant narrower essential classification, so that some disjuncts would be eliminated;

let me omit this complication henceforth.) That alleged event does indeed deserve rejection. Twice over: both because it is disjunctive and because the disjuncts are excessively extrinsic. But the perfectly good event which is the loss of heat by the poker is specified as such *accidentally*, not essentially. It is essentially a decreasing of molecular motion. This event does not occur in otherworldly regions where the poker loses heat by losing caloric fluid. But it does occur in regions where the molecular motion decreases and yet the poker does not lose heat, these being regions of worlds where something besides molecular motion is what occupies the heat-role. In those worlds, of course, the event fails to fits its thisworldly accidental description as a loss of heat.

Likewise in general. Whenever some term nonrigidly designates the occupant of a role, and that role could be occupied in a variety of ways, the term becomes unsuitable for essential specification of events. If being fragile means having some or another basis for a disposition to break when struck, and if many different properties could serve as such bases (under this- or otherworldly laws), then no genuine event is essentially classifiable as the window's being fragile. There is a genuine event which is accidentally classifiable in terms of fragility; essentially, however, it is a possession of such-and-such molecular structure, that being the actual basis of the window's fragility. (This event is an unchange, but I haven't rejected those.) I think this observation gives the sense in which, as Prior *et al.* say, dispositions are "inefficacious."[15] And if I am right to think that mental states are definable as occupants of causal roles, then no genuine event is essentially classifiable as my being in pain. There are pain events, no doubt of it; but they are pain events only accidentally, just as pain itself is a property that only contingently occupies its role and deserves its name. Essentially, the events are firings of neurons, perhaps—unless "firing" and "neuron" also are terms for occupants of roles, in which case we must get more physical before we finally reach an essential classification.

If there are no extrinsic or disjunctive events to be caused, still there are extrinsic or disjunctive truths about regions to be explained. They can be explained, of course. And their explanations can be mostly or entirely causal, even if my theses about causal explanation of events do not apply directly. The explanandum truth is made true by a pattern of

genuine, occurrent events. (This making true is logical, not causal.) These events have their causal histories. Explanatory information about the explanandum truth consists in part of noncausal information about the truth-making pattern itself: what sort of pattern it is, and what events comprise it. And it consists in part of information about the causal histories of the events that comprise the pattern. As usual, explaining means providing some explanatory information. The serving provided may consist of noncausal information about the pattern, or causal information about its events, or some of each. Why did Xanthippe become a widow? Because she was married to Socrates at the time of his death. (Noncausal.) Because Socrates was made to drink hemlock. (Causal, with the noncausal background most likely presupposed.) Why did Fred talk or walk then? Because he talked (noncausal) and he did that because he had just heard a joke he couldn't keep to himself (causal).