

from Lewis D. (1986) *Philosophical papers*, vol. II, Oxford University Press

Postscripts to "Causation"

A. Piecemeal Causation

Suppose that c and e are large, prolonged processes, each composed of many smaller events. Suppose it is not true (or not clearly true) that e , taken as a whole, causally depends on c , taken as a whole; suppose even that they are not connected by a chain of causal dependence. It may nevertheless be that c and e are divisible into parts in such a way that every part of e is causally dependent on (or connected by a chain of causal dependence to) some part of c . In that case we might well simply speak of c as a cause of e , though it is not so under the analysis I gave.

Self-sustaining processes exhibit piecemeal causation. For instance, suppose a public address system is turned up until it howls from feedback. The howling, from start to finish, is an event. If it had not occurred, it would not have occurred; but this is certainly not counterfactual dependence between *distinct* events, therefore it does not qualify as causal dependence on my account. Nor is there a closed causal loop, as in time travel stories, in which the howling causes itself because it depends causally on some distinct event which in turn depends causally on it. So it is not true, on my account, that the howling *taken as a whole* causes itself. What is true is that the howling

end p.172

causes itself piecemeal. It is divisible into parts in such a way that each part except the first is caused by an earlier part, and each part except the last causes a later part. This causing of part by part is unproblematic: cause and effect are distinct events, wherefore their counterfactual dependence qualifies as causal. We might well say that the howling causes itself; this is to be accepted, but only in a derivative sense. Similarly, if two prolonged events sustain one another,

each causes the other piecemeal. The example of the howling illustrates this case also: the sound in the air sustains the signal in the wires, and *vice versa*.

It may be that when we speak of causation in history we are often speaking of piecemeal causation.¹ A depression causes a wave of bankruptcies: what are we to make of this? If the depression had not occurred. . . . That is puzzling. To suppose away an entire depression takes us a long way from actuality. And the farther we depart from actuality, the more we lose control over our counterfactuals. For the more different respects of similarity and difference we have to balance, the more of a problem it is that we have left it vague just how to do the balancing, so the less clearly we know what is and what isn't to be held fixed in our counterfactualizing. (For instance, what if many of the firms that went broke came into existence during, and because of, the depression itself? Shall we hold their existence fixed in asking what would have happened without the depression?) But the depression is a big event that is divisible into many parts. Although it is hard to say what would have happened without the entire depression, it is comparatively easy to say that without this or that event which was part of the depression, this or that one of the bankruptcies would not have taken place. Now, our counterfactuals are much more under control, because they stay much closer to home. So even if it is unclear what the depression taken as a whole might have caused, it is at any rate clear that various parts of it caused the various bankruptcies. That is to say that the depression was at least a piecemeal cause of the wave of bankruptcies.

There is a well-known dilemma about actions. Consider an action of raising my arm. First something goes on within my brain; then signals go out my nerves; then my muscles contract; and as they do, my arm rises. There seems to be a conflict between two things we want to say. (1) The action of raising my arm is a prolonged event with diverse parts. It is the whole causal process just described. It may begin within

end p.173

me, but it is not over until my arm rises. Its earlier parts cause its later parts, and its final part is the bodily movement. But (2) just as my action of raising a flag would be an event that causes a flag to rise, so my action of raising my arm is an event that causes my arm to rise. The raising, whether of flag or of arm, is so-called because it causes the rising.

Distinguish the inclusion of one event in another from mere involvement of one in the naming of the other. So far as involvement in naming goes, the two cases are on a par. For the flag and the arm alike, the raising deserves its name only if, and perhaps only after, a rising ensues. But with respect to inclusion, the two cases seem to differ. If I raise a flag by delayed action, I can be done raising it long before it rises. My action is over when I have done my part; the process that ends when the flag is up consists of more than just my action. (Beware ambiguity: the phrase "my raising the flag" might denote just my action, or it might denote the whole affair.) But if I raise my arm by delayed action—say that I have *very* sluggish nerves—then it takes me a long time to raise my arm. In this case, my part of the process is the whole of the process. So long as the signal is traveling through my sluggish nerves, so long as my muscles are contracting and my arm is rising, my part of the affair is still going on.

I would like to assent to both (1) and (2); the apparent obstacle is that we have two events, the raising and the rising, and according to (1) they are not wholly distinct, yet according to (2) one causes the other. But if this is a case of piecemeal causation, we have no problem. If an early part of the raising causes the rising which is a late part of the raising, we may still say simply that the raising causes the rising; just as, when an early part of the depression causes a bankruptcy which is a later part of the depression, we may still say simply that the depression causes the bankruptcy.

There is a second version of the problem. The rising of my arm is not the only event which is caused by the initial inner part of my action and yet takes place before my arm has risen. The same may be true of various side effects, events which definitely are not to be included as parts of the action. Suppose, for instance, that the nerves leading into my arm are monitored so that whenever I raise my arm the nerve signal produces a trace on an oscillograph. Because I can produce the trace by raising my arm, we ought to be free to say that my action causes the trace. And

yet the trace appears before the arm rises. Shall we say that the effect precedes its cause? Or that the action which causes the effect is over sooner than we think? Neither: it is a case of piecemeal

end p.174

causation. Like the rising of the arm, the trace on the oscillograph is caused by an initial part of the action, and thereby is caused by the action.²

B. Chancy Causation

In the paper, I confined my discussion to the deterministic case for the sake of brevity.³ But I certainly do not think that causation requires determinism. (Hence I regard "causality" as a naughty word, since it is ambiguous between "causation" and "determinism.") Events that happen by chance may nevertheless be caused. Indeed, it seems likely that most actual causation is of just this sort. Whether that is so or not, plenty of people do think that our world is chancy; and chancy enough so that most things that happen had some chance, immediately before-hand, of not happening. These people are seldom observed to deny commonplace causal statements, except perhaps when they philosophize. An analysis that imputes widespread error is *prima facie* implausible. Moreover, it is dishonest to accept it, if you yourself persist in the "error" when you leave the philosophy room. We had better provide for causation under indeterminism, causation of events for which prior conditions were not lawfully sufficient.

One kind of chancy causation is already covered by my analysis, with no modification needed: *c* occurs, *e* has some chance of occurring, as it happens *e* does occur; but if *c* had not occurred, then *e* would have had no chance at all of occurring, and so would not have occurred. Then *e* depends causally on *c*, and *c* is a cause of *e*, according to my original analysis. So far, so good.

(Some would object to my step from "*e* would have had no chance of occurring" to "*e* would not have occurred." They say that things

end p.175

with no chance at all of occurring, that is with probability zero, do nevertheless happen; for instance when a fair spinner stops at one angle instead of another, yet any precise angle has probability zero. I think these people are making a rounding error: they fail to distinguish zero chance from infinitesimal chance. Zero chance is *no* chance, and nothing with zero chance ever happens. The spinner's chance of stopping exactly where it did was not zero; it was infinitesimal, and infinitesimal chance is still *some* chance.)

But there is a second case to be considered: *c* occurs, *e* has some chance *x* of occurring, and as it happens *e* does occur; if *c* had not occurred, *e* would still have had some chance *y* of occurring, but only a very slight chance since *y* would have been very much less than *x*. We cannot quite say that without the cause, the effect would not have occurred; but we can say that without the cause, the effect would have been very much less probable than it actually was. In this case also, I think we should say that *e* depends causally on *c*, and that *c* is a cause of *e*.

It does not matter whether *x* itself, the actual chance of the effect, is high or low. Suppose you mischievously hook up a bomb to a randomizer—a genuinely chancy one, if need be one that works by counting clicks in a counter near a radioactive source. If you set the randomizer to a high probability, that makes it likely that your act of setting up the bomb will cause an explosion. If you set the randomizer to a low probability, that makes it less likely that your act will cause an explosion. But no matter how you set the randomizer, if the bomb does chance to go off, then your act does cause the explosion. For no matter how you set the randomizer, we can be sure that the explosion would have been very much less probable still if you hadn't set up the bomb at all.

(You took it in stride when you read my words: if you set the randomizer low, that makes it less likely that your act will cause an explosion. That proves my point. For suppose that improbable events cannot be caused: the actual chance *x* has to be high, or at least has to exceed some lowish threshold, in order to have a case of causation. Then if you set the randomizer low enough, that doesn't just make it *unlikely* that your act will cause an explosion—it makes it downright

impossible. But "unlikely" did *seem* the right word. "Don't worry—set the randomizer below 0.17% and you can't *possibly* cause an explosion."—Not so!)

Several points of clarification may be helpful. (1) Chances are time-dependent: an event may have different chances at different times

end p.176

before it occurs. The actual chance x of e is to be its chance at the time immediately after c ; and the counterfactual is to concern chance at that same time. (2) I do not assume that there is some y that would definitely have been the chance of e in the absence of c . Maybe so, maybe not. Maybe in that case the chance of e might have had any of various values, all of them much less than the actual chance x . In saying that without c , e would have had some chance y much less than x , "some chance y " is a quantifier whose scope is limited to the consequent of the conditional. (3) "Much less" means less by a large factor—not by a large difference. If x is already small, the difference of y and x could not be large. It is x that sets the standard for how small the chance of e must be without c . We could have one case in which the absence of a cause would lower the chance of an effect from 100% to 10%; another in which the lowering would be from 10% to 1%; yet another in which the lowering would be from 1% to 0.1%; . . . and all would count equally as cases of chancy causal dependence. So it will not do to simplify our counterfactual and say that without c , the chance of e would be low *simpliciter*. (4) A chance event may be caused; but we should not say that it is caused to happen *rather than not*. Contrastive causal statements differ from plain ones. According to what I say about contrastive questions and statements in "Causal Explanation" (in this volume), there can be no contrastive causal explanation of why a chance event occurs rather than not.

Many probabilistic theories of causation share the motivating idea that a cause increases the probability of the effect. Mine differs from some of the others in two respects.⁴ First, it is meant to apply to causation in the single case: causation by one particular event of another event, not conduciveness of one kind of event to another kind. Hence its probabilities are single-case chances, as opposed to finite or

end p.177

limiting frequencies. You may not like single-case chances—I don't either—but I cannot see how to make sense of certain well-established scientific theories without them. If we need them anyway, we may as well use them here. (I discuss single-case chances, and the reason for disliking them, elsewhere in this volume; see "A Subjectivist's Guide to Objective Chance," especially the final section and Postscript C; also my discussion of Humean supervenience in the introduction.)

Second, my analysis is in terms of counterfactual conditionals about probability; not in terms of conditional probabilities. If we try to use an inequality of conditional probabilities to express that event c raises the probability of event e , we run into a well-known difficulty. The inequality may well hold not because c causes e , but rather because c and e are two effects of a common cause. One cure is to use fancier conditional probabilities: conditionalize not just on the absence of c , but on that together with a specification of background. Then the problem is to say, preferably without circular mention of causation, what information should be included in this background.

But even if that problem can be solved, another remains. Conditional probabilities, as standardly understood, are quotients. They go undefined if the denominator is zero. If we want to say, using conditional probabilities, that c raises the probability of e , we will need probabilities conditional on the non-occurrence of c (plus background, perhaps). But there is no guarantee that this conditional probability will be defined. What if the probability that c occurs (given background) is one? What if c has been predetermined through all of past time—what if its probability has *always* been one, so that even by going back in time we cannot find a non-zero chance of c 's failing to occur? For that matter, what if we want to apply our probabilistic analysis of causation to a deterministic world in which all probabilities (at all times) are extreme: one for all events that do occur, zero for all that don't? The requisite conditional probabilities will go undefined, and the theory will fall silent. That is not acceptable. Earlier, I said that it would not do to impute error to

indeterminists who accept commonplace causal statements; therefore we cannot accept an analysis of causation that works only under determinism. Likewise it would not do to impute error to determinists who accept commonplace causal statements; therefore we cannot accept an analysis that works only under indeterminism. An adequate analysis must be neutral. It must work in both cases. And it must work in a uniform way, for it does not seem that our concept of causation is disjunctive. A probabilistic analysis (of single-case causation) that uses conditional probabilities is

end p.178

not neutral. It is made for indeterminism. My analysis, on the other hand, can serve alike under indeterminism or determinism.⁵

My motivating idea is that causes make their effects more probable; but that is written into the analysis of causal dependence, not of causation itself. As in my original analysis, we have causation when we have a causal chain: one or more steps of causal dependence. The effect need not depend on the cause directly. When we have causation without direct causal dependence, as in some cases of preemption, it is not necessarily true that the cause at the beginning of the chain raises the probability of the effect at the end. The cause might lower the probability of the effect, or might leave it unchanged. At each step in the chain, we have a cause raising the probability of its immediate effect. But since counterfactuals are not transitive, that does not settle whether there is raising over the entire chain.

Suppose we have two redundant systems to produce the same effect. One is much more reliable than the other—that is, much less subject to random failure part way along the causal chain. The reliable system is already started; left to itself, it will very probably produce the effect. But I do not leave it to itself. There is a switch that both turns off the reliable system and turns on the unreliable system, and I throw this switch. As luck would have it, the unreliable system works. The effect ensues, just as it would probably have done without my act. My act did not make the effect more probable, but rather less, since I put the unreliable system in place of the reliable one. Nevertheless, I did cause the effect. And the reason is plain if we consider some intermediate event in the causal chain that actually occurred, something that happened well after the reliable system was already turned off. That event was part of the working of the unreliable system, so it would not have occurred, or at least it would have been improbable, if I had not thrown the switch. But by the time of the intermediate event, the reliable system was already out of action. So without that event, the effect would not have occurred, or at least it would have been very improbable. (Here it is crucial that the counterfactual be governed by a similarity relation that does not conduce to backtracking; see "Counterfactual

end p.179

Dependence and Time's Arrow" in this volume.) My act raised the probability of the intermediate event, and thereby caused it. And the intermediate event raised the probability of the effect, and thereby caused it. So my act caused a cause of the effect, and thereby caused it—despite lowering its probability.⁶

I have said that if distinct events *c* and *e* both occur, and if the actual chance of *e* (at a time *t* immediately after *c*) is sufficiently greater than the counterfactual chance of *e* without *c*, that implies outright that *c* is a cause of *e*. Some philosophers find this counterintuitive. They would correct me thus.

No; if there would have been some residual chance of *e* even without *c*, then the raising of probability only makes it *probable* that in this case *c* is a cause of *e*. Suppose, for instance, that the actual chance of *e*, with *c*, was 88%; but that without *c*, there would still have been a 3% probability of *e*. Then most likely (probability 97%) this is a case in which *e* would not have happened without *c*; then *c* is indeed a cause of *e*. But this just might be (probability 3%) a case in which *e* would have happened anyway; then *c* is not a cause of *e*. We can't tell for sure which kind of case this is.

It is granted, *ex hypothesi*, that it would have been a matter of chance whether *e* occurred. Even so, the objection presupposes that the case must be of one kind or the other: either *e* definitely *would* have occurred without *c*, or it definitely would *not* have occurred. If that were so, then indeed it would be sensible to say that we have causation only in case *e* definitely would not have occurred

without *c*. My original analysis would serve, the amendment suggested in this postscript would be unwise, and instead of having a plain case of probabilistic causation we would have a probable case of plain causation.

But I reject the presupposition that there are two different ways the world could be, giving us one definite counterfactual or the other. That presupposition is a metaphysical burden quite out of proportion to its intuitive appeal; what is more, its intuitive appeal can be explained away.

The presupposition is that there is some hidden feature which may or may not be present in our actual world, and which if present would

end p.180

make true the counterfactual that *e* would have occurred anyway without *c*. If this counterfactual works as others do, then the only way this hidden feature could make the counterfactual true is by carrying over to the counterfactual situation and there being part of a set of conditions jointly sufficient for *e*.

What sort of set of conditions? We think at once that the set might consist in part of laws of nature, and in part of matters of historical fact prior to the time *t*, which would together predetermine *e*. But *e cannot* be predetermined in the counterfactual situation. For it is supposed to be a matter of chance, in the counterfactual situation as in actuality, whether *e* occurs. That is stipulated as a hypothesis of the case. When an event is predetermined, there cannot be any genuine chance that it will not happen. Genuine chance gives us the residue of uncertainty that is left after *all* laws and prior conditions have been taken into account.

(Here I assume that we are not dealing with an extraordinary situation, involving time travel perhaps, in which the normal asymmetries of time break down, and the past contains news from the future. That is fair. The objection concerns what should be said about *ordinary* cases of probabilistic causation.)

So the hidden feature must be something else. But what else can it be? Not the historical facts prior to *t*, not the chances, not the laws of nature or the history-to-chance conditionals that say how those chances depend on the prior historical facts. For all those are already taken account of, and they suffice only for a chance and not a certainty of *e*.

There is the rest of history: everything that happens after *t*. These future historical facts are not relevant to the chances at *t*; *e* can still have a chance of not occurring even if there are facts of later history that suffice for its occurrence. As there will be: if it does occur, that is itself a fact of later history. In the terminology of "A Subjectivist's Guide to Objective Chance" (in this volume) later history is "inadmissible." So perhaps that is where the hidden feature of the world is to be found.

But this also will not do. For we know very well that if we give weight to future similarities, so that facts of the later history of our world tend to carry over into counterfactual situations, then we will get into trouble. We will get counterfactuals that seem false in themselves, and that also yield false conclusions about causation. We must make sure, either by fiat or else by tailoring our standards of similarity to exploit the *de facto* asymmetries of time, that future similarities will

end p.181

normally carry no weight. (See "Counterfactual Dependence and Time's Arrow" in this volume.) Features of our actual future history may be well hidden, sure enough, and they might well enter into sets of conditions and laws sufficient to postdetermine *e*, but what they will not do is carry over into the counterfactual situation without *c*.

(Normally. I am forced to admit exceptions of two kinds, for reasons discussed in Postscript D to "Counterfactual Dependence and Time's Arrow" in this volume. If a reconvergence to actual history could be accomplished without widespread miracles or quasi-miraculous coincidences, then I would admit that actual future history carries over into the counterfactual situation; and I would admit that the absence of such quasi-miracles carries over. But I think the first cannot apply to the truth of counterfactuals at a world like ours, and the second could apply only to the special case where *e* itself would be quasi-miraculous. So these exceptions are not relevant to our present discussion.)

So the hidden feature must be something else still: not a feature of the history of this world, and also not a feature of its chances, or of the laws or conditionals whereby its chances depend on its history. It fails to supervene on those features of the world on which, so far as we know, all else supervenes. To accept any such mysterious extra feature of the world is a serious matter. We need some reason much more weighty than the isolated intuition on which my opponent relies.⁷ Without such a reason, it would be better to suppress the intuition.

end p.182

Which is all the easier if it rests on a mistake in the first place; and I think it does. I suspect that my opponent is someone who has not wholeheartedly accepted my stipulation of the case in question. Stipulation or no, he remains at least somewhat inclined to think that the case involves not genuine chance, but a kind of counterfeit chance that is compatible with determinism. (See Postscript B to "A Subjectivist's Guide to Objective Chance" in this volume.) Perhaps he clear-headedly thinks that counterfeit chance is all the chance there could ever be, and so is all that could be meant by the word "chance." Or perhaps he thinks double, and thinks of the case half one way and half the other.

If it *is* a case of counterfeit chance, then his objection is well taken. For then *e* is after all predetermined one way or the other, both in actuality and in the counterfactual situation without *c*; but predetermined partly by details of prior historical fact that are far too minute to be discovered in advance. So we do indeed have an unproblematic hidden feature of the actual world—namely, the relevant configuration of minute details—that carries over to the counterfactual situation and there joins in predetermining the outcome one way or the other.

That is all very well, but then his objection is off target. I was not speaking of a case of counterfeit chance, I insist, but of a different case: probabilistic causation of a genuine chance event. If my opponent believes that my case is impossible because counterfeit chance is all the chance there can be, let him say so; but let him not reinterpret my case to fit his own doctrines.

When my opponent says that either *e* would have occurred without *c* or else *e* would not have occurred without *c*, he sounds like Robert Stalnaker.⁸ But his position is not the same, though he accepts the same disjunction of counterfactuals, and Stalnaker's defense of such disjunctions is no use to him. My opponent thinks there are two relevant ways the world might be; one of them would make true one of the disjointed counterfactuals, the other would make true the other, so the disjunction is true either way. Stalnaker, like me, thinks there is only one relevant way for the world to be, and it does not make either counterfactual determinately true. But Stalnaker, unlike me, thinks the disjointed counterfactuals are true or false relative to alternative arbitrary resolutions of a semantic indeterminacy; what makes the counterfactuals

end p.183

lack determinate truth is that different resolutions go different ways; but every resolution makes one or the other true, so the disjunction is determinately true despite the complementary indeterminacies of its disjuncts. A resolution of an alleged semantic indeterminacy is not a hidden fact about the world; and that is the difference between Stalnaker and my opponent. Stalnaker disagrees with me on a small point of semantics; my opponent, on a large point of ontology. A resolution of an indeterminacy might indeed be *mistaken* for a hidden fact about the world—Stalnaker suggests, plausibly, that such mistakes are common. So if we accepted Stalnaker's view on the point of semantics, that would give us a second way to explain away my opponent's problematic intuition.

C. Insensitive Causation

Killing, so they say, is causing to die. I am sure that I—and likewise you, and each of us—have caused ever so many people to die, most of them people yet unborn. Acts of mine are connected to their deaths by long chains of causal dependence.⁹ But I have never killed anyone—I hope.

For instance, suppose I write a strong recommendation that lands someone a job; so someone else misses out on that job and takes another; which displaces a third job-seeker; this third job-seeker goes elsewhere, and there meets and marries someone; their offspring and all their descendants forevermore would never have lived at all, and *a fortiori* would never have died, and so presumably their deaths would not have occurred, but for my act.¹⁰ Maybe there is a time after which *every* death that occurs is one that would not have occurred but for my act. It would be strange to single out my act as *the* cause of all those deaths. But it is *a* cause of them, under my analysis and also according to our common usage. And still I deny that I have ever killed.

For a still more striking case, consider the Big Bang. This event, I

end p.184

take it, is a cause of *every* later event without exception. Then it is a cause of every death. But the Big Bang did not kill anyone.

So killing must be a special kind of causing to die. But what distinguishes this special kind of causation?

Not that there must be one single step of causal dependence, as opposed to an intransitive chain. An act of killing can be a preempting cause. It can be you who kills the victim, even though another killer was standing by who would have done the job for you—causing the victim to die the very same death—if he had not seen you lay the poison yourself.

Not that the chance of the effect must be high. If you hook up a bomb to a randomizer and hide it in a crowded place, and it happens to go off, you can kill no matter how low you set the chance.

Not that the causal chain must be short. You can kill by delayed action. If you set a hidden time bomb with a thousand-year fuse, you may well kill someone yet unborn.

Not that the chain must be simple. You can kill someone by means of a lethal Rube Goldberg machine.

Not that the chain must be foreseeable. You can kill someone no matter how good your reasons were for thinking the gun was not loaded; or no matter how unfeasible it would have been for you to discover in advance his lethal allergy to what you fed him.

Not that the chain must pass through no later human actions.¹¹ If you kill by setting a baited mantrap, or by making a gift of poisoned chocolates, your unsuspecting victim's action is an intermediate step in the causal chain whereby you kill him. In other cases, an action by a third party may be an intermediate step: you make a gift of poisoned chocolates to the host, who offers them to the guest.

Perhaps a cluster of these conditions, inadequate if taken one by one, would work to distinguish the kind of causing that can be killing. I think not. But the counterexamples get too contrived to be very persuasive: imagine a lethal Rube Goldberg machine with a randomizer at one step, a thousand-year fuse at another, an alternative waiting in reserve at another, dependence on some action of the unsuspecting victim at another, and no way to discover how it works.

end p.185

I suggest a different way to distinguish the right kind of causing: by its insensitivity to circumstances. When an effect depends counterfactually on a cause, in general it will depend on much else as well. If the cause had occurred but other circumstances had been different, the effect would not have occurred. To the extent that this is so, the dependence is sensitive. Likewise if a causal chain consists of several steps of causal dependence, we can say that the chain is sensitive to the extent that its steps are. (On average? Or at worst?) Sensitivity is a matter of degree, however. It may be that the causation depends on an exceptionally large and miscellaneous bundle of circumstances all being just right. If any little thing had been different, that cause would not have caused that effect. But sometimes causation is comparatively insensitive to small differences in the circumstances. When my strong recommendation causes lives and then deaths, that is comparatively sensitive causation—there are many differences that would have deflected the chain of events. But if you shoot at your victim point-blank, only some very remarkable difference in circumstances would prevent his death. The same is true if you set a Rube Goldberg machine, or a delayed-action

bomb, working inexorably toward its lethal outcome. The case of the bomb with a randomizer also is comparatively insensitive: the bomb might very well have chanced not to go off, but it isn't the fine details of the circumstances that would make the difference.

Jonathan Bennett restates my suggestion this way: killing requires "that the causal chain run through a stable and durable structure rather than depending on intervening coincidental events."¹² A lethal Rube Goldberg machine may work in many steps, it may be full of thousand-year fuses and randomizers and alternatives waiting in reserve, its working may require the responses of unsuspecting agents, there may be no way to discover how it is built or understand how it would work; and yet it may be no less "stable and durable" for all that, and the causal chain running through it may be far more independent of "intervening coincidental events" than are most of the causal chains in the wider world.

So it seems that the reason why a lot of causing to die is not killing is, at least partly, that the causing to die in killing must be causation of a comparatively insensitive kind. And if this is so for killing, perhaps it is so likewise for other causatives. Consider the ways in which you can and can't make, break, wake, or bake things.

end p.186

Insensitivity is not the same thing as any of the unsatisfactory conditions that I considered above; but of course it is connected to several of them. *Ceteris paribus*, shortness and simplicity of the chain will make for insensitivity; insensitivity, in turn, will make for foreseeability. The more the chain depends on a lot of circumstances being just right, the harder it is for a would-be predictor to know all he needs to know about the circumstances. The sensitivity of the chain is an obstacle to prediction. Unforeseeability does not imply sensitivity, since any of many other obstacles to prediction might be at work. But unforeseeability sets a minimum standard. If a chain is insensitive enough that you can predict it, then it is insensitive enough that you can kill by it. Perhaps our common knowledge of what can normally be predicted sets a common standard for everyone. Or perhaps the standard varies. What if you are much better than I am at predicting chains that are somewhat sensitive? I am inclined to say that if so, then indeed you can kill in ways that I cannot. If your act and mine cause death by chains that are exactly alike, and if the duplicate chains are insensitive enough to fall within your powers of prediction but sensitive enough to frustrate mine, then you kill but I do not.¹³

My suggestion faces a problem. Recall that you can kill by a causal chain that has someone else's action as an intermediate step: you give someone poisoned chocolates, he unsuspectingly serves them to his guest, and thereby you kill the guest. (It is true as well that the host unwittingly kills the guest. But that is beside the point; the question is whether you kill the guest, and I submit that you and the host both do.) But if you tell the host that the chocolates are poisoned, and you order or hire or coerce or persuade him to serve them anyway, then it seems that you do not kill the guest. You may be no less guilty, morally and in the eyes of the law, than if you had killed him; or no less praiseworthy, if the guest was Hitler. Be that as it may, it seems that you don't kill by getting someone else to kill *knowingly*. Why not, on the suggestion I have advanced? It seems that if someone else is ready to kill knowingly when ordered or hired or coerced or persuaded to, his readiness well might be a stable and durable structure; so that by depending on this readiness, the causal chain from your action to a death well might be fairly insensitive to fortuitous circumstances.

I reply that indeed that might be so, and nevertheless we might speak as if it were not so. That would be no surprise. Part of our habitual

end p.187

respect for other people consists in thinking that they are sensitive to a great variety of considerations, and therefore not easy to predict or control. It is all very well to take for granted that someone is ready to offer a guest what he takes to be harmless chocolates; to that extent, it is not disrespectful to regard his dispositions as a stable and durable structure. Offering chocolates is no big deal. It is another thing to take for granted that someone can be ordered or hired or coerced or persuaded to kill knowingly. That is to take him altogether too much for granted. The relevant disrespect lies not in thinking him willing to kill; whether that is disrespectful depends on the

circumstances and the victim. Rather, it lies in thinking of his readiness to kill as stable and durable, inexorable, insensitive to fortuitous circumstances of the case, so that he is disposed to make weighty choices with unseemly ease.

Such disrespect might be well deserved. We might know very well that this dull thug before us would never think twice about killing for a small fee. Therefore, we might be sure that when you hire him, the causal chain from your action to the victim's death is as inexorable and insensitive as if it had passed instead through some strong and sturdy machine. But we might know this, and yet be halfhearted in putting our mouths where our minds are. Some vestige of our habitual respect might well influence how we speak. If I am right, when you cause death by hiring this thug, you are in literal truth a killer, no less than the thug himself is. If we deny it, I suggest that we are paying the thug a gesture of respect—insincere, undeserved, yet unsurprising.

That was an uncompromising version of my reply. I can offer an alternative version that runs as follows. If you hire the thug just considered, you are *not* in literal truth a killer. The truth conditions for "kill" are not just a matter of insensitive causation. They make an exception for insensitive causal chains that run through someone else's action of knowingly killing. However, insensitivity remains the underlying idea. The extra twist in the truth conditions is not just a brute complication of the concept; it is there, understandably, thanks to our respectful presumption that a causal chain through someone else's weighty decision will not be insensitive. The two versions agree about what we say, and why we say it; they differ only about what is literally true. *Ceteris paribus* it is bad to claim that we say what we know is literally false; but *ceteris paribus* it is bad to build complicating exceptions into the conditions of literal truth. Between the version that does one and the version that does the other, I think there is little to choose. I am not even confident that there is a genuine issue between the two.

end p.188

D. Causation by Omission¹⁴

An omission consists of the nonoccurrence of any event of a certain sort. To suppose away the omission is, exactly, to suppose that some event of the given sort does occur. We say that omissions may be caused, and may cause, and I have no wish to deny this. I would like to be able to provide for causation by omission within the general framework given in this paper and in "Events" (in this volume). Unfortunately, I do not see how to make it fit with all that I say in general about events and about their causal dependence. So, one way or another, a special case it must be.

Omissions as effects are no special problem. I must allow in any case that sometimes, by causing suitable events, causes can create a pattern of events; and that a fact can supervene on this pattern even when there is no genuine event that can be called the obtaining of that fact; in which case the causes of the events in the pattern can also be said to cause that fact to obtain.¹⁵ For instance, it is at least a fact that Xanthippe became a widow. I think there is no genuine event that can be called Xanthippe's becoming a widow. But the causes of her marriage together with the causes of Socrates's death may nevertheless be said to have caused her to become a widow: they caused genuine events that comprised a pattern on which the fact that she became a widow supervened. Certainly this fact is not beyond the reach of causal explanation. Likewise I can say that various distractions caused Fred to omit the precautions he should have taken; and in saying this, I needn't grant that there was any such thing as an event of omission. If there are events of omission, well and good. But I don't need them as effects.

Do I need them as causes? There are two opposite strategies that I might follow, and a third which is a compromise between those two. One way or another, all of them treat causation by omission as a special case. While I would guess that any of the three could be made to work, I am not in a position to prove it by presenting fully developed versions. I am not sure how much the three really differ; certainly some of their difference is just terminological.

end p.189

The first strategy accepts that there are events of omission. What is more, there are events essentially specifiable as omissions. For instance, Fred's omission of precautions, essentially

specifiable as such, is an event that would have occurred no matter how he omitted them, no matter what else he did instead; and that could not have occurred if he had taken the precautions. For any event, there are necessary and sufficient conditions, normally hard to state, for that very event to occur. Some descriptions of an event are built into its conditions of occurrence; others are not. The first strategy says that the description of this event as an omission *is* built in. Then to suppose counterfactually that this event of omission does not occur is equivalent to supposing that Fred does take the precautions. So the counterfactual analysis of causation can apply to events of omission just as it does to all other events; and it is safe to say, as we ought to, that the effects which depend causally on Fred's omission are those which would not have occurred if he had taken the precautions. This strategy requires no exception to what I say about causation in general.

But it does require an exception to what I say about events in general. For I say that a theory of events, if it is built to serve the needs of my analysis of causation, must reject overly disjunctive events. An alleged event would be disjunctive if, or to the extent that, it could have occurred in various dissimilar ways. (The point is not that its conditions of occurrence could be formulated as a disjunction—anything can be formulated as a disjunction—but that they could be formulated as a disjunction of overly varied disjuncts.) An alleged event that is essentially specifiable as a talking-or-walking, and which could have occurred either as a talking or as a walking, is an example of what ought to be rejected. The reason is that if it were accepted as an event, then it could qualify as a cause; but it is intuitively very wrong to say that the talking-or-walking causes anything. But if we are to accept events of omission, in the way we are considering, then we may not reject disjunctive events without exception. For an event of omission, essentially specifiable as such, is highly disjunctive. Fred omits the precautions if he does something else during the period in which he was supposed to attend to them. So there are as many different ways for the event of omission to occur as there are alternative ways for Fred to spend the time. An event essentially specifiable as an omission amounts to an event essentially specifiable as a sleeping-or-loafing-or-chatting-or- . . . with a disjunct for everything Fred might do other than attending to the precautions. If omissions are accepted as genuine events and as causes, while other alleged disjunctive events are rejected,

end p.190

that makes causation by omission a special case.¹⁶ The unfinished business for this strategy, of course, is to draw the line: how do we distinguish the genuine omissions from other alleged events that we should still reject? For instance, I think we ought not to say that the laws of nature, or other regularities, cause things; yet regularities may be made out to be omissions on a cosmic scale—the universe omits to contain events that would violate them. We must somehow deny that we have here a genuine event of omission.

The second, opposite, strategy says that there are no events of omission. Then there is no need to make a place for them within a theory of events, and no need to worry that they would be like other alleged events that are to be rejected. So far, so good. But in that case, I need to make an exception to what I say about causation itself. For it is not to be denied that there is causation by omission; and I cannot analyze this in my usual way, in terms of counterfactual dependence between distinct events. Instead I have to switch to a different kind of counterfactual for the special case. The counterfactual is not: if event *c* (the

end p.191

omission) had *not* occurred. . . . It is rather: if some event of kind *K* (the omitted kind) *had* occurred. . . .

But if we use special counterfactuals for the special case, that opens several questions. Again we need to draw a line. I thought it necessary to block some counterexamples against a counterfactual analysis of causation by insisting that counterfactual dependence was to be between events. If we give that up, what new line shall we retreat to? As before, alleged causation by the laws of nature, regarded as cosmic omissions, will illustrate the problem. Also, I thought it necessary to insist on distinctness between events that stand in causal dependence, and by distinctness I meant more than nonidentity. (See "Events.") But how does distinctness apply to

causation by omission? Fred sleeps, thereby omitting precautions against fire and also omitting precautions against burglary. Two distinct omissions?

The third, compromise strategy accepts events of omission as causes; but this time, the events of omission are not essentially specified as such. Fred omits the precautions, sleeping through the time when he was supposed to attend to them. His nap was a genuine event; it is not objectionably disjunctive. There are many and varied ways in which he could have omitted the precautions, but there is just one way that he did omit them. We could plausibly say, then, that his nap *was* his omission of precautions. But accidentally so. His nap could have occurred without being an omission of precautions: if (1) that very nap had been taken somewhat later, with the precautions seen to beforehand; or conceivably (2) if he had taken the precautions somehow in his sleep, or (3) if that very nap could have been taken by someone else, or (4) if the precautions had not been his responsibility. (I take it that (2)-(4) are problematic in various ways; so I rest my case mainly on (1).) And an omission of precautions might very well have occurred without being that nap: he might have stayed awake and done any of many other things instead of attending to the precautions. Still, *as it was*, the nap was what happened instead of the taking of precautions. So we may call it an event of omission, though we do not thereby capture its essence. We can have events of omission, so understood, and still reject disjunctive events without exception.

But this third strategy, like the second, demands special counterfactuals for the special case. Even if Fred's nap was his omission of precautions, it is one thing to suppose that this very event did not occur, and it is another thing to suppose that no event that occurred (this or any other) was an omission of the precautions. It is one thing to suppose away the event *simpliciter*, another thing to suppose it away *qua*

end p.192

omission. It is the second counterfactual supposition, not the first, that is relevant to causation by omission. For it is the second supposition that is equivalent to supposing that the precautions were taken. But this is special. In other cases the relevant counterfactuals are those that suppose away an event *simpliciter*; we do not in general need to suppose away events *qua* satisfying some or other accidental specification.

As with the second strategy, resort to special counterfactuals for the special case threatens to undo our defenses against various counter-examples. It remains to be seen how, if at all, those defenses could be rebuilt. This strategy, like the others, leaves us with unfinished business.

E. Redundant Causation¹⁷

Suppose we have two events c_1 and c_2 , and another event e distinct from both of them; and in actuality all three occur; and if either one of c_1 and c_2 had occurred without the other, then also e would have occurred; but if neither c_1 nor c_2 had occurred, then e would not have occurred. Then I shall say that c_1 and c_2 are *redundant causes* of e .

(There might be redundant causation with a set of more than two redundant causes. There might be probabilistic redundant causation, in which e would have had some small chance of occurring even if neither c_1 nor c_2 had occurred. There might be stepwise redundant causation without direct dependence, as described by Louis Loeb.¹⁸ I pass over these complications and consider redundant causation in its simplest form.)

As in my definition of ordinary causation, the counterfactuals concern particular events, not event-kinds. So it is not redundant causation if you shoot a terminal cancer patient—or, for that matter, a healthy young mortal—who would sooner or later have died anyway. Without your act

end p.193

he would have died a different death: numerically different, because very different in time and manner. The particular event which is the death he actually dies would not have occurred. If you shoot a man who is being stalked by seven other gunmen, that *may* be redundant causation—the answer depends partly on details of the underdescribed case, partly on unsettled standards of how much difference it takes to make a different event. If you shoot a man who is simultaneously

being shot by seven other members of your firing squad, that doubtless is redundant causation. The exact number of bullets through the heart matters little.

If one event is a redundant cause of another, then is it a cause *simpliciter*? Sometimes yes, it seems; sometimes no; and sometimes it is not clear one way or the other. When common sense delivers a firm and uncontroversial answer about a not-too-far-fetched case, theory had better agree. If an analysis of causation does not deliver the common-sense answer, that is bad trouble. But when common sense falls into indecision or controversy, or when it is reasonable to suspect that far-fetched cases are being judged by false analogy to commonplace ones, then theory may safely say what it likes. Such cases can be left as spoils to the victor, in D. M. Armstrong's phrase. We can reasonably accept as true whatever answer comes from the analysis that does best on the clearer cases. It would be still better, however, if theory itself went indecisive about the hard cases. If an analysis says that the answer for some hard case depends on underdescribed details, or on the resolution of some sort of vagueness,¹⁹ that would explain nicely why common sense comes out indecisive.

In my paper, I distinguished one kind of case—preemption with chains of dependence—in which common sense delivers clear positive and negative answers, and my counterfactual analysis succeeds in agreeing. I left all other cases of redundant causation as spoils to the victor, doubting that common-sense opinions about them would be firm and uncontroversial enough to afford useful tests of the analysis.

end p.194

Now I would distinguish more varieties of redundant causation. Sometimes my analysis, as it stands, agrees with clear common-sense answers, positive or negative. Sometimes it reproduces common-sense indecision. Sometimes I am still content to leave far-fetched cases as spoils to the victor. But sometimes it seems that additions to my original analysis are needed.

I consider first a class of cases distinguished by doubt as to whether they exhibit redundant causation at all. I have already mentioned one example: you shoot a man who is being stalked by seven other gunmen. As it actually happens, the man dies on Tuesday morning, face down on the ground, his heart pierced by your bullet, with an entry wound in his back and an exit wound in his chest. Without your act he would have died on Wednesday evening, slumped in a chair, his heart pierced by someone else's bullet, with an entry wound in his chest and an exit wound in his back.


Is it that without your act he would have died a different death—numerically different because somewhat different in time and manner?²⁰ If so, there is no redundancy. The particular death he actually died depends counterfactually on your act, without which that very

end p.195

event would not have occurred. This is straightforward causation. Or is it rather that without your act he would have died the very same death—numerically the same, despite slight differences in time and manner? If so, there is genuine redundancy. In that case your act would be a redundant cause; whether it would be a cause *simpliciter* awaits our discussion of the varieties of genuine redundant causation.

It is hard to say which is true. It would remain hard, I think, no matter how fully we described the details of what actually happened, and of what would have happened under our counterfactual hypothesis.

Here is another example. Suppose three neurons are hooked up thus.



Suppose that a neuron fires if stimulated by the firing of one or more other neurons connected to it by a stimulatory synapse (shown by a forward arrowhead). But suppose—fictitiously, I believe—that a neuron fires much more vigorously if it is doubly stimulated than if it is singly stimulated. Neurons C_1 and C_2 fire simultaneously, thereby doubly stimulating E , which

fires vigorously. Is this vigorous firing of **E** a different event from the feeble firing that would have occurred if either one of C_1 and C_2 had fired alone? Then we have joint causation, in which the effect depends counterfactually on each of the causes, and there is no redundancy. Or is it that numerically the same firing would have occurred, despite a difference in manner, with single stimulation? Then we have redundant causation. Again it is hard to say, and again the difficulty cannot be blamed on underdescription of the details.

Call an event *fragile* if, or to the extent that, it could not have occurred at a different time, or in a different manner. A fragile event has a rich essence; it has stringent conditions of occurrence. In both our examples we have redundant causation if the effect is not too fragile, ordinary causal dependence on joint causes otherwise.

Don't say: here we have *the events*—how fragile are *they*? Instead it should be: here we have various candidates, some more fragile and some less—which ones do we call the events? (For instance under my proposal in "Events," in this volume, the candidates will be smaller and larger classes of possible spatiotemporal regions, more and less

end p.196

tightly unified by similarity.) Properly posed, the question need not have a fully determinate answer, settled once and for all. Our standards of fragility might be both vague and shift.

As of course they are. You can say: the performance should have been postponed until the singer was over his laryngitis; then *it* would have been better. You can just as well say, and mean nothing different: the performance should have been cancelled, and another, which would have been better, scheduled later to replace it. There's no right answer to the question how fragile the performance is. Not because there is something—the performance—with an indeterminate size in logical space! But because there are various things, with various sizes, and we haven't troubled to decide which one is "the performance." Likewise every region of the earth has exact boundaries and a determinate size. Silicon Valley, whatever exactly that is, is no exception. However we haven't decided exactly how big a region is called "Silicon Valley." That's why there's no right answer to the question whether these words (written on the Stanford campus) were written in Silicon Valley.

So there may be no right answer to the question whether we have a case of joint causation without redundancy; or whether instead we have a case of redundant causation, which might or might not count as causation according to considerations to be discussed later. The answer depends on the resolution of vague standards of fragility. If common sense falls into indecision and controversy over such cases, that is only to be expected.

It is a common suggestion to adopt extreme standards of fragility, and thereby make away with redundant causation altogether. Even if a man is shot dead by a firing squad, presumably it would have made *some* minute difference to the time and manner of his death if there had been seven bullets instead of eight. So if you fired one of the eight bullets, that made some difference; so if his death is taken to be very fragile indeed, then it would not have occurred without your act. Under sufficiently extreme standards of fragility, the redundancy vanishes. Even this turns out to be a case in which the effect depends on each of several joint causes. Likewise for other stock examples of redundancy.

(Suppose we did follow this strategy wherever we could. Wouldn't we still have residual cases of redundancy, in which it makes *absolutely* no difference to the effect whether both of the redundant causes occur or only one? Maybe so; but probably those residual cases would be mere possibilities, far-fetched and contrary to the ways of this world. Then we could happily leave them as spoils to the victor. For we could plausibly suggest that common sense is misled: its habits of thought are

end p.197

formed by a world where every little thing that happens spreads its little traces far and wide, and nothing that happens thereafter is quite the same as it would have been after a different past.)

Extreme standards of fragility would not fit a lot of our explicit talk about events. We do say—within limits!—that an event could have been postponed and could have happened differently.

But this is not a decisive objection. The standards that apply within the analysis of causation might differ from those that apply in explicit talk.

What matters more is that extreme standards would not fit a lot of our negative judgements about causation itself. Extreme fragility of effects would make for spurious causal dependence in many quite ordinary cases. It would make more trouble than it cures.²¹

For instance, suppose there was a gentle soldier on the firing squad, and he did not shoot. If the minute difference made by eight bullets instead of seven is enough to make a different event, then so is the minute difference made by eight instead of nine. So if the victim's death is so very fragile that it would not have occurred without your act, equally it is so fragile that it would not have occurred without the gentle soldier's omission. If by reason of fragility the death depends causally on your act, then equally it depends causally on the omission. So the gentle soldier caused the death by *not* shooting, quite as much as you caused it by shooting! This is a *reductio*.

That case may puzzle us because it involves at least an appearance of redundancy, and also because it involves causation by omission. But the problem arises for cases without these complications. Boddie eats a big dinner, and then the poisoned chocolates. Poison taken on a full stomach passes more slowly into the blood, which slightly affects the time and manner of the death. If the death is extremely fragile, then one of its causes is the eating of the dinner. Not so.

To be sure, resolution of vagueness is influenced by context; and I can imagine a special context in which we might after all agree that the eating is a cause of the death. Pleased that Boddie is dead but horrified that the death was lingering, the poisoner says: if only he hadn't eaten, *this* wouldn't have happened—and by "this" he means the death, taken as very fragile. Maybe indeed that context makes it right to say that the eating caused the death. But it is also right, certainly in other contexts and probably even in this one, to say what is true under more lenient

end p.198

and more ordinary standards of fragility: namely, that the eating did *not* cause the death.²²

So if we wanted to make away with the stock examples of redundant causation, what we would require is not a uniformly stringent standard of fragility, but rather a double standard—extremely stringent when we were trying to show that an effect really depends on its alleged redundant causes, but much more lenient when we were trying to agree with common-sense judgements that an effect is not caused by just anything that slightly affects its time and manner. It is not out of the question that there should be such a double standard. But if there is, an adequate theory of causation really ought to say how it works. (The changes of standard noted above, brought on by contextual pressures, are not the ones we want—they cut across cases with and without apparent redundancy.) To say how the double standard works may not be a hopeless project, but for the present it is not so much unfinished as unbegun.

Extreme fragility of effects might get rid of all but some far-fetched cases of redundant causation, but it leads to trouble that we don't know how to control. Moderate fragility gets rid of some cases and casts doubt on others, but plenty are left. Our topic has not disappeared.

So I return now to genuine redundant causation, including the doubtful cases when taken under standards of fragility that make them genuine. I divide it into *preemption* and (*symmetrical*) *overdetermination*.²³ In a case of preemption, the redundant causes are not on a par. It seems clear that one of them, the *preempting cause*, does the causing; while the other, the *preempted alternative*, waits in reserve. The alternative is not a cause; though it could and would have been one, if it had not been preempted. There is the beginning of a causal process running from the preempted alternative to the effect. But this process does not go to completion. One effect of the preempting cause is to cut it off. In a case of overdetermination, on the other hand, there is no such asymmetry between the redundant causes. It may or may not be

end p.199


clear whether either is a cause; but it is clear at least that their claims are equal. There is nothing to choose between them. Both or neither must count as causes.

First, preemption. It is clear what answer we want—the preempting cause is a cause, the preempted alternative is not—and any analysis that does not yield that answer is in bad trouble. It is

easy for me to say why the preempted alternative is not a cause: the effect does not depend on it. My problem is to say why the preempting cause *is* a cause, when the effect does not depend on it either. (A regularity analysis of causation has the opposite problem: why is the preempted alternative not a cause, when it is part of a set of conditions jointly sufficient for the effect?)

I subdivide preemption into *early* and *late*. In early preemption, the process running from the preempted alternative is cut off well before the main process running from the preempting cause has gone to completion. Then somewhere along that main process, not too early and not too late, we can find an intermediate event to complete a causal chain in two steps from the preempting cause to the final effect. The effect depends on the intermediate, which depends in turn on the preempting cause. (Or, in cases with more than one preempted alternative, we might need more steps.) We have a causal chain of stepwise dependence between the cause and the effect, even if not dependence *simpliciter*; and since causation is transitive, we take the ancestral of dependence. Thus I say that *c* is a cause of *e* if there is a sequence *c*, . . . , *e* of events, consisting of *c* and *e* and zero or more intermediates, with each event in the sequence except the first depending on the one before. (Normally all these events would be distinct, and in temporal order; but I do not require this. See Postscript F, below.)

This is the variety of preemption that I discussed in the paper. To illustrate it, let us have another system of neurons.



Besides stimulatory synapses from one neuron to another, as before, we now have an inhibitory synapse as well (shown by a backward arrowhead). A neuron normally fires if stimulated, but not if it is inhibited

end p.200

at the same time. Neurons *C*₁ and *C*₂ fire; thereby starting two processes of firing which make their separate ways toward neuron *E*. The main process, which begins with the firing of *C*₁, goes to completion. But the alternative process, which begins with the firing of *C*₂, is cut short: because neuron *I* is inhibited, the neurons shown dotted never fire. There is also a branch process, diverging from the main process. The junction event where it diverges is the firing of neuron *J*. It is this branch process that cuts off the alternative process by inhibiting neuron *I*. The main and alternative processes—the one actual, the other partly unactualized—merge with the firing of neuron *M*; and proceed thence to the final effect, the firing of neuron *E*.

Thus the firing of *C*₂ is the preempted alternative. It is not a cause of the firing of *E* because there is no direct dependence, and neither is there any stepwise dependence via an intermediate. The firing of *C*₁ is the preempting cause. The firing of *D* is our intermediate event. It depends counterfactually on the firing of *C*₁; the firing of *E* depends on it; and thereby we have our two-step chain of dependence from the preempting cause to the effect. For by the time of the firing of *D*, the alternative process was already doomed. The alternative process was doomed as soon as neuron *J* fired; though it was not yet cut off, the branch process that was going to cut it off had already diverged from the main process. So if the firing of *D* had not occurred, *both* processes would have failed, and the firing of *E* also would not have occurred.

Don't say that if *D* had not fired, that would mean that it had not been stimulated, and that would mean that the neurons to its left on the main process would not have fired, and so neuron *I* would not have been inhibited, and so the alternative process would have gone to completion and *E* would have fired after all. That is backtracking; and backtracking counterfactuals, however legitimate in other contexts, are out of place in tracing causal dependence. (See "Counterfactual Dependence and Time's Arrow" in this volume.) Of course it is not just to deal with early

preemption that we must avoid backtracking; as is explained in the paper, the avoidance of backtracking is needed also to solve the problems of effects and of epiphenomena.


We have some choice which event goes along the main process to take as our intermediate. The firing of J comes too early: the effect does not depend on it, since without it the alternative process would not have been cut off. The firing of M comes too late: it lies on the unactualized alternative process as well as on the main process, and so does not depend on the preempting cause. But anything in between would do. What makes the solution possible is that there exists some

end p.201

intermediate event in the gap between too early and too late. And so it is, generally, in cases of early preemption. Thus we distinguish the genuine cause from its preempted alternative, as we should, even though either one by itself would have sufficed to cause the effect.²⁴

Late preemption is harder. Our solution cannot succeed unless there is a sufficient gap between too early and too late; if not-too-early is already too late, there is no place for an intermediate event to complete a chain of stepwise dependence.

There are two far-fetched ways in which this problem might arise. The first way involves action at a temporal distance. Suppose that in our previous example, we remove all the neurons between J (too early) and M (too late).



In their place, suppose we have some law of delayed action that directly connects the firings of J and M. Iff J fires, then M fires a certain time later (as in the original example) but that is absolutely all there is to it—there is no connection between the two neurons, and no continuous causal process between their two firings. That is possible, I take it, though it goes against what we take to be the ways of this world. In such a case, we have no intermediate event to complete our chain of dependence.

The second way involves infinite multiple preemption. We have infinitely many preempted alternatives, and infinitely many cut-off alternative processes. Suppose for simplicity that the main process and its unactualized alternatives merge only at the final effect. (Otherwise the problem would be the same, but with the point of merging in place of the final effect.) Then any other event on the main process is not too late to depend on earlier events along that process. The problem is to

end p.202

find an intermediate event that is not too early to take the penultimate place in our chain of stepwise dependence—that is, to find an event on which the final effect depends. Such an event has to come late enough that by the time it occurs, all of the infinitely many alternative processes are doomed. Any one of the alternative processes is eventually doomed, so there is an event that comes late enough so far as it is concerned. Likewise for any finite set. But since there are infinitely many alternatives, there may be no event before the final effect that comes after *all* the alternative processes are doomed. Suppose one of them is doomed 128 seconds before the final effect, another only 64 seconds before, another only 32 seconds before, . . . Then at no time before the final effect are all of them doomed. Then there is no intermediate event on which the final effect depends. Our causal chains of stepwise dependence can get as close as we like to the final effect, but they never can reach it. Then there is no stepwise dependence between the effect and what seems to be its preempting cause.²⁵


I do not worry about either of these far-fetched cases. They both go against what we take to be the ways of this world; they violate the presuppositions of our habits of thought; it would be no surprise if our common-sense judgements about them therefore went astray—spoils to the victor! Common sense does judge them to be cases of causal preemption, in which what seems to be a preempting cause is indeed a cause, despite the lack of either direct or stepwise dependence. But an

analysis that disagrees may nevertheless be accepted. It would be better to agree with common sense about these cases, to be sure, but that is not an urgent goal.

Unfortunately there is another variety of late preemption, quite commonplace and not at all far-fetched; and there it *is* an urgent goal to agree with common sense. Again we have what seems to be a preempting cause, hence a cause *simpliciter*, but no dependence and no stepwise dependence. Here my analysis seems to be in trouble. These are cases in which an alternative process is doomed only when the final effect itself occurs. The alternative is cut off not by a branch process that diverges from the main process at a junction event before the effect is reached, but rather by a continuation of the main process beyond the effect. Shooting a man stalked by seven other gunmen would be a case of this kind, if it is a case of redundant causation at all, and if the other gunmen desist only when they see him dead. Another case would be

end p.203

this system of neurons. Again we start with the simultaneous firings of neurons C_1 and C_2 , which redundantly cause the firing of E.



I ignored such cases when I wrote the paper, and for many years afterward. My reason must have been that there is a ready-made solution: fragility of the effect. If the alternative process is only doomed by the effect itself, and if at the time of the effect it is not yet complete, then the alternative process must run more slowly. So if it had been left to produce the effect, the effect would have been delayed. Without the firing of neuron C_1 (the seeming main cause) the firing of neuron E would have been delayed by the time it takes for three extra neurons to fire; if you had not shot the man on Tuesday morning, he would not have died until Wednesday evening; and so on, for all such cases. We can devise cases in which the delay is very short, but we can never get rid of it altogether. (Or not without resort to instantaneous or backward causation. But then the case becomes far-fetched, not worrisome, spoils to the victor.) If the effect is taken to be fragile, then the delay would suffice to give us a numerically different event instead of the effect that actually occurred. We would have causal dependence without redundancy, thus agreeing with common sense that your shooting the man on Tuesday, or the firing of C_1 , or whatever, is indeed a cause.

But my reason for ignoring these cases was a bad reason, because the ready-made solution is a bad solution. Fragility of the effect is no better as a remedy for these cases of late preemption than it is as a remedy for redundant causation generally. To deal with all the cases, including those where the delay is very short and there is not much difference in manner to go with it, we need extreme standards of fragility; uniformly extreme standards are no good because they will give us lots of spurious causal dependence; so we need a double standard; and that might be workable, for all we know, but we don't know how to make it work. There are two problems. One is that a double standard must be principled. We need some definite rule to tell us when we should raise the standard: when is dependence among fragile versions relevant,

end p.204

and when is it not, to causation among the original robust events? The second problem is that a stringent standard may give the wrong answer. Let c_1 be a preempting cause of e , and let c_2 be the preempted alternative, in a case of late preemption. Without c_1 , e would have been delayed; and so a more fragile version of e would not have occurred at all. So far, so good. But it may also be that some side effect of c_2 substantially influences the time and manner of e ; in which case, unfortunately, a version of e that is fragile enough to depend on c_1 may depend on c_2 as well. Indeed, it may take more fragility to give us the dependence on c_1 that yields the right answer than it does to give us the dependence on c_2 that yields the wrong answer. Though I don't reject the

fragility approach out of hand, I don't see how to make it work.²⁶ So I am inclined to prefer a different solution, though it is more of a departure from my original analysis in the paper.

Leaving the problem of late preemption in abeyance, consider this question. Suppose we have processes—courses of events, which may or may not be causally connected—going on in two distinct spatiotemporal regions, regions of the same or of different possible worlds. Disregarding the surroundings of the two regions, and disregarding any irrelevant events that may be occurring in either region without being part of the process in question, what goes on in the two regions is exactly alike. Suppose further that the laws of nature that govern the two regions are exactly the same. Then can it be that we have a causal process in one of the regions but not the other? It seems not. Intuitively, whether the process going on in a region is causal depends only on the intrinsic character of the process itself, and on the relevant laws. The surroundings, and even other events in the region, are irrelevant. Maybe the laws of nature are relevant without being intrinsic to the region (if some sort of regularity theory of lawhood is true) but nothing else is.

Intuitions of what is intrinsic are to be mistrusted, I think. They too often get in the way of otherwise satisfactory philosophical theories. Nevertheless, there is some slight presumption in favor of respecting them. Let us see where this one leads us.

A process in a region may exhibit a pattern of counterfactual dependence that makes it causal, according to my original analysis. Its later parts may depend counterfactually on its earlier parts (later and earlier in time, normally, but all I require is that there be dependence with

end p.205

respect to some order); and in particular, its last event may depend on its first. (We will provide for stepwise dependence later.)

Now suppose that some process in some region does not itself exhibit this pattern of dependence; but suppose that in its intrinsic character it is just like processes in other regions (of the same world, or other worlds with the same laws) situated in various surroundings. And suppose that among these processes in other regions, the great majority—as measured by variety of the surroundings—do exhibit the proper pattern of dependence. This means that the intrinsic character of the given process is right, and the laws are right, for the proper pattern of dependence—if only the surroundings were different, and different in any of many ways. According to my original analysis, this process is nevertheless not causal. Thanks to its special bad surroundings, it is a mere imitation of genuine causal processes elsewhere. But that goes against our motivating intuition.

So we might extend the analysis. Suppose that there exists some actually occurring process of the kind just described, and that two distinct events *c* and *e* are the first and last in that process. Then let us say that *e* *quasi-depends* on *c*. We might wish to count that as one kind of causation, based derivatively on counterfactual dependence even though there is no dependence between those two events themselves. As before, we must take an ancestral to ensure that causation will come out transitive; thereby providing not only for chains of stepwise dependence, but also for chains of stepwise quasi-dependence, or mixed chains. To this end we could redefine a *causal chain* as a sequence of two or more events, with either dependence or quasi-dependence at each step. And as always, one event is a *cause* of another iff there is a causal chain from one to the other.

That would solve the problem of late preemption, both in the commonplace cases that worry me and in the far-fetched cases that do not. For the problem is that we seem to have a causal process starting with a preempting cause, and ending with the final effect; and yet this process does not exhibit the proper pattern of counterfactual dependence, not even if we count stepwise dependence. Segments of it do exhibit dependence, but we cannot patch these segments together to make a chain that reaches all the way to the effect. What spoils the dependence is something extraneous: the presence alongside the main process of one or more preempted alternatives. Without them, all would be well. Hold fixed the laws but change the surroundings, in any of many ways, and we would have the dependence that my original analysis requires for causation. But as is, we have quasi-dependence instead of

end p.206

dependence. So if we extend the analysis, and allow causation by quasi-dependence, that solves our problem. We then can agree with common sense that we have genuine preemption, and genuine causation by the preempting cause.²⁷

The extended analysis, which allows causation by quasi-dependence, is more complicated than my original analysis; and it is less purely a counterfactual analysis, though of course counterfactual dependence still plays a central role. The complication would be objectionable if it were just a hoky gimmick to deal with late preemption, but it is not just that. For what it is worth, we also have independent motivation in the intuition of intrinsicness. While I would still welcome a different solution to the problem of late preemption, within my original analysis, I now think that the extended analysis may well be preferable.


This completes my discussion of preemption. I now turn to the other variety of redundant causation: overdetermination, with nothing to break the symmetry between the redundant causes. When I wrote

end p.207

the paper, I thought that all such cases were alike; that a counterfactual analysis would inevitably deny that the redundant causes in overdetermination are causes *simpliciter*; and that it did not matter much what the analysis said, since all such cases were spoils to the victor for lack of firm common-sense judgements.

All that is wrong. An important paper by Martin Bunzl changes the picture greatly.²⁸ Bunzl observes that when we examine stock examples of overdetermination in detail, we can very often find an intermediate event—call it a *Bunzl event*—that satisfies two conditions. First, the Bunzl event is jointly caused, without redundancy, by the same events that are redundant causes of the final effect. Second, the Bunzl event seems clearly to be a cause (often a preempting cause) of the final effect. Cases of overdetermination are not all alike, because there are different kinds of Bunzl events (at least three) and also because there are some possible cases, far-fetched perhaps, with no Bunzl events at all. A counterfactual analysis does not deny that the redundant causes are causes *simpliciter* of the final effect, provided it can agree that they are causes of a Bunzl event and that the Bunzl event in turn is a cause of the effect. The cases should not all be left as spoils to the victor, because once a Bunzl event is noticed, it becomes clear to common sense that we have genuine causation.

One kind of Bunzl event is a preempting cause in a case of late preemption. This system of neurons illustrates it. Here B is an especially lethargic neuron. It will not fire if singly stimulated, but it will if doubly stimulated.



As usual, the simultaneous firings of C_1 and C_2 are redundant causes of the firing of E. But also they are joint causes, without redundancy, of a Bunzl event: namely, the firing of B. And that is a preempting cause of the final effect. The preemption is late: the two alternative processes,


end p.208

those that run from the firings of C_1 and C_2 taken separately, are cut off only because of the effect itself. It is the firing of E itself that inhibits neurons I_1 and I_2 . We must apply whatever solution we favor for late preemption generally. If somehow we had a double standard of fragility, we might say that the firing of E is extremely fragile, and would not have occurred (though E would still have fired) without the firing of B. Or, probably better, we might say that despite a lack of direct or stepwise dependence, we have causation by quasi-dependence. Thanks to the intrinsic character of the course of events running from the firing of B to the firing of E, we would have had counterfactual dependence if just such a course of events had occurred under the same laws but in

any of various different surroundings. According to the extended analysis, that means that the firing of B is a cause of the firing of E. Either way, we say as we should that the firing of B causes the firing of E; and therefore, by transitivity, the firing of C_1 and C_2 are both causes of E.

Not only is the firing of B a preempting cause; we can also think of the entire course of events as a case of *self*-preemption. The firings of C_1 and C_2 , our redundant causes, both preempt and are preempted. Taken together as joint causes, they preempt themselves taken separately.

A second kind of Bunzl event is a preempting cause in a case of early preemption, as in the following system of neurons.




Again the simultaneous firings of C_1 and C_2 are redundant causes of the firing of E, and joint causes without redundancy of the firing of the lethargic neuron B. As in simpler cases of early preemption, the firing of D completes a chain of stepwise dependence: it depends on the firing of B, and in turn the firing of E depends on the firing of D. So there is a three-step causal chain of dependence from the firing of C_1 to the firing

end p.209

of E; and likewise from the firing of C_2 . More simply, there is a two-step chain, since the firing of D also depends directly on the firing of C_1 , and likewise on the firing of C_2 . The firing of B is a Bunzl event; so is the firing of D; and so are various other intermediate events on the chain. Again we have self-preemption by our redundant causes: the firings of C_1 and C_2 taken jointly preempt themselves, taken separately.

This looks complicated. But just the same sort of early self-preemption can happen in much simpler cases of overdetermination, as follows.

The third kind of Bunzl event is a fragile intermediate. Earlier, we considered a case of fragility of the effect, involving a neuron that would fire vigorously if doubly stimulated, feebly if singly stimulated. We considered that under moderate and reasonable standards of fragility, hence without any problematic double standard, we might say that the vigorous firing and the feeble firing would differ enough in manner to make them numerically different events. If we place the fragile vigorous firing as an effect, what we have is not redundant causation at all. But if we place it as an intermediate, it can be the Bunzl event in a case of overdetermination. Here is such a case, with B as the neuron that may fire either vigorously or feebly.



The vigorous firing of B that actually occurs depends on both of the simultaneous firings of C_1 and C_2 . Without either one of these causes it would not have occurred. The feeble firing of B that would have occurred with only one of them would not have been the same event. But also the firing of E depends on the firing of B. So each of our redundant causes is connected to the final effect by a two-step causal chain of dependence. Not by direct dependence: if only one of C_1 and C_2 had fired, so that B fired feebly, E would still have been stimulated and its firing would have been

very little different. This is not a case that can be treated by fragility of the effect, or not under moderate standards of fragility.

(My solution depends on assuming that if the intermediate event—the vigorous firing of B—had not occurred, then B would not have fired at all. It isn't that the vigorous firing would have been replaced by

end p.210

a feeble firing, differing only just enough not to be numerically the same. That may seem to go against a similarity theory of counterfactuals—wouldn't the minimal change to get rid of an event be one that replaces it with a barely different event? Not so; a similarity theory needn't suppose that just any sort of similarity we can think of has nonzero weight. It is fair to discover the appropriate standards of similarity from the counterfactuals they make true, rather than *vice versa*. (See "Counterfactual Dependence and Time's Arrow" in this volume.) And we certainly do not want counterfactuals saying that if a certain event had not occurred, a barely different event would have taken its place. They sound false; and they would make trouble for a counterfactual analysis of causation not just here, but quite generally.)

The case looks simpler than the self-preemption cases above; but it is really much the same. The process from the redundant causes jointly through the vigorous firing to the effect goes to completion. The two alternative processes from the redundant causes taken singly through the feeble firing to the effect are cut short when the feeble firing does not occur. The feeble firing is prevented by the double stimulation of **B**, and that is an event in the main process.

Still there is one important difference from previous cases. When we have a fragile intermediate, as opposed to the sorts of Bunzl events considered above, there is room for serious indeterminacy. Just as our vague and shifty standards of fragility may leave it unsettled whether we have a fragile effect, so they may leave it unsettled whether we have a fragile intermediate. Then they may leave it unsettled whether we have overdetermination with or without a Bunzl event. If that is what decides whether the redundant causes are causes *simpliciter*, that question too may have no right answer.

So I turn to the last variety of redundant causation: overdetermination without a Bunzl event, including doubtful cases when taken under standards of fragility that give no relevant fragility either in the effect or in the intermediates. According to my original analysis, the redundant causes in such a case are not causes *simpliciter*, because there is neither direct nor stepwise dependence. But the extended analysis would disagree. There is quasi-dependence of the effect on each of the two redundant causes, and if we allowed causation by quasi-dependence, that would make the redundant causes count as genuine causes of the effect.

Also, the original analysis will say that in cases where it is doubtful whether there is a fragile effect or intermediate, then it is likewise doubtful whether the redundant causes are causes *simpliciter*. Whereas

end p.211

the extended analysis would say that in such cases the redundant causes are causes, though the reason why is left doubtful. The first analysis would be better suited to explain indecision and controversy, the second would be better suited to explain positive judgements.

I used to think that all cases of overdetermination, as opposed to preemption, could be left as spoils to the victor; and that is what I still think about these residual cases. All the more so, given Bunzl's discussion of what we find when we look at realistic cases in microscopic detail, without simplifying idealizations. For it seems that cases without Bunzl events require phenomena with perfectly sharp thresholds, whereas thresholds under the laws of this world are imperfectly sharp. Thus I am content to say that these cases may go one way or the other. The decision will depend on what strategy emerges as victor in the cases that really matter—namely, the commonplace cases of late preemption.

I should dispel one worry: that if we ever decline to count redundant causes as genuine causes, then we will be left with gaps in our causal histories—no cause at all, at the time when the redundant causes occur, for a redundantly caused event. That is not a problem. For consider the

larger event composed of the two redundant causes. (I mean their mereological sum. *Not* their disjunction—I do not know how a genuine event could be the disjunction of two events both of which actually occur. It would have to occur in any region where either disjunct occurs. Hence it would have to occur twice over in one world, which a particular event cannot do. See "Events" in this volume.) Whether or not the redundant causes themselves are genuine causes, this larger event will be there to cause the effect. For without it—if it were completely absent, with neither of its parts still present, and not replaced by some barely different event—the effect would not occur. For *ex hypothesi* the effect would not occur if both redundant causes were absent, and to suppose away both of them is just the same as to suppose away the larger event that is composed of them.

F. Self-Causation

My requirement that cause and effect be distinct applies to causal dependence, but not to causation generally. Two events are distinct if they have nothing in common: they are not identical, neither is a proper part of the other, nor do they have any common part. Despite the truth of the appropriate counterfactuals, no event depends causally

end p.212

on itself; or on any other event from which it is not distinct. However, I do allow that an event may cause itself by way of a two-step chain of causal dependence: *c* depends on *d* which depends in turn on *c*, where *d* and *c* are distinct. Likewise for longer closed causal loops; or for loops that lead from an event back not to itself but to another event from which it is not distinct. Thus I have taken care not to rule out the sort of self-causation which appears in time-travel stories that I take to be possible. (See "The Paradoxes of Time Travel" in this volume.)

But no event can be self-caused unless it is caused by some event distinct from it. Indeed, no event can be caused at all unless it is caused by some event distinct from it. Likewise no event can cause anything unless it causes some event distinct from it.

Suppose we think of the entire history of the world as one big event. It is not caused by any event distinct from it; else that distinct event both would and would not be part of the entire history. Likewise it does not cause any event distinct from it. So it has no causes or effects at all. Not as a whole, anyway. Its parts, of course, do all the causing there is in the world.

Some philosophers wish to believe only in entities that have some causal efficacy.²⁹ Either they must reject such totalities as the big event which is the whole of history, or else they should correct their principle. They might admit those inefficacious things that could have been efficacious if, for instance, there had been more of history than there actually was. Or, more simply, they might admit those inefficacious things that are composed entirely of efficacious parts.

end p.213