

This is a close-to-final draft of a paper that appeared in The Journal of Philosophy, April 2000. To cite, please refer to the published article.

¹Aspect Causation

L.A. Paul

Yale University

While skiing, Suzy falls and breaks her right wrist. The next day, she writes a philosophy paper. Her right wrist is broken, so she writes her paper using her left hand. (Assume, as seems plausible, that she isn't dexterous enough to write it any other way, e.g., with her right foot.) She writes the paper, sends it off to a journal, and it is subsequently published. Is Suzy's accident a cause of the publication of the paper?²

Of course not. Below, I will show that none of the major contenders for a theory of events coupled with a theory of causation succeeds against examples like that of Suzy's accident, and that the reason for this derives from an underlying tension between our beliefs about events and our goals for theories of causation. I will then argue that property instances should be taken, in the first instance, as the causal relata, and propose an analysis of causation that I call *aspect causation*.

Aspect causation combines elements of regularity theories and David Lewis's new influence theory of causation³ with property instances. Combining lawful entailment with influence handles problems involving redundant causation for regularity and counterfactual based theories. Changing the causal relata to property instances resolves transitivity problems and allows us to develop an account of events without holding our theory of causation hostage to our theory of

event individuation.

1. Event Causation: Problems

First, we need some background and terminology: actual events can be defined according to a continuum from *fine grained* to *coarse grained*. Speaking loosely, the less fine grained a theory of individuation specifies events to be, the fewer events the theory implies. There is a parallel continuum from *fragile* to *robust* for the individuation of events with respect to modal contexts, needed for counterfactual theories of causation. The less fragile the event, the weaker the requirements events in other possible worlds must meet to be numerically the same as the event that actually occurred.

One of two approaches is usually used to analyze causation: a regularity approach or a counterfactual approach. Roughly, versions of regularity or covering-law analyses for events state that an event c causes an event e iff the occurrence of c , together with the right conditions and right regularities or laws, is sufficient for the occurrence of e . It is a consequence that causation is transitive: if c is a cause of d , and d is a cause of e , then c is a cause of e . Counterfactual analyses in their simplest form hold that for any two actual, distinct events c and e , e depends causally on c iff e depends counterfactually on c , i.e., had c not occurred then e would not have occurred. Causation is usually distinguished from causal dependence in order to ensure transitivity: c is a cause of e iff there is a chain of causal dependencies running from c to e . Both accounts require transitivity to avoid important counterexamples,⁴ and both require an acceptable theory of eventhood in order to provide an acceptable (reductive) analysis of causation. In order to provide an underpinning for analyses of causation based on events as causal relata, several

characterizations of the identity conditions for events have been put forward.

Fragile events Under a fragility view, one must hold that when I pick up my cup of coffee, had I picked up the coffee a millisecond later, had the coffee swirled in a slightly different fashion, or had it been a fraction of a degree hotter, the event of my picking up the coffee cup would—necessarily—have been a different event. Since anything that brings about a new event counts as a cause of it, the view gives us a plethora of spurious causes. If an explosion on the sun makes the summer day on which I drink my coffee a fraction of a degree hotter, thus keeping my coffee ever so slightly warmer than it would have been had the explosion not occurred, that explosion counts as a cause of my drinking my coffee.

If events are fragile, then Suzy's skiing accident is among the causes of her writing her paper with her left hand, and her writing of the paper is a cause of it being published. For fragility theorists, writing the paper with her left hand is necessarily a different event than writing the paper in some other way. So the skiing accident is a cause of the event that actually occurred: the event *e* of the paper's-being-written-with-the-left-hand, since the skiing accident is part of a sufficient condition for *e*, and if the skiing accident had not occurred, *e* would not have occurred. Given that the skiing accident occurred, if the writing of the paper occurs at all, it occurs with the left hand, since after the accident she can't write it any other way. Under a counterfactual analysis that ensures transitivity and treats events as fragile, if the event of the writing of the paper with the left hand had not occurred, then the publication would not have occurred. Such a fragility theorist is committed to the view that Suzy's skiing accident is a cause of her paper's being published.

The implausibility of such conclusions increases with every link in the causal chain, and confronts advocates of fragility with a multitude of counterintuitive cases. Perhaps the fragile event theorist can provide some sort of an account that can make one-step cases such as the sun's explosion causing my drinking of my coffee seem less counterintuitive. But the implausibility of holding, e.g., that the skiing accident is among the causes of Suzy publishing her paper, and the ease with which such two-step (or multi-step) examples can be constructed shows us that explaining one-step cases is not good enough.

Fine grained events. Using fine grained events as part of a theory of causation gives us many of the same undesirable results as the use of fragile events, coupled with a need to accept an implausible ontology. Jaegwon Kim's view is perhaps the most well known: he combines a regularity approach with a very fine grained theory of events. For Kim, events are identified by their constitutive triples, comprised of an individual, a property exemplified by that individual, and the time when the property is exemplified. Although an event may exemplify many properties, it is the constitutive triple which, in a sense, defines the event.⁵ For Kim, we have two numerically different events when they differ in their constitutive triples: i.e., when they differ in their constitutive individuals, properties or times.

When cat C. Louise sneezes loudly, according to Kim at least two events occur in that region of spacetime: C. Louise's sneezing and C. Louise's sneezing loudly. Moreover, when she sneezes, she sneezes in the kitchen, on a summer morning while dislodging a flea; so we also have the simultaneously occurring events of a sneezing on a summer morning, a sneezing in the kitchen, and a sneezing while dislodging a flea. We can continue our modification of the constitutive

property indefinitely, so we have an infinite number of events occurring in the region of C. Louise's sneezing. For Kim, these events are all different but not necessarily distinct: in some sense (although he does not develop the notion), he thinks the many events that occur when C. Louise sneezes are all 'included' in her sneezing.⁶ The claim is supposed to be similar to the claim that C. Louise is made up of many different C. Louises, one for each hair or molecule that C. Louise might lose but still remain C. Louise. C. Louise is made up of many spatiotemporal parts, and many of these parts are less maximal than all of C. Louise, but these proper parts are enough to count as C. Louise themselves.

But the attempt to run an analogy between included events and spatiotemporal parts fails.⁷ The difference between events that are qualitatively richer than other events (such as the difference between a sneezing in the kitchen, and a sneezing at 10 am) is not purely analogous to the difference between spatiotemporal parts. If Kim were claiming that the event of C. Louise's sneezing in the kitchen and sneezing at 10 am occupied slightly different spatiotemporal regions, then we could see how one event could overlap another, and how part of one event could include part of another event. But both of these events are said to occupy the very same spatiotemporal region. The only way to make sense of the 'parts' analogy is to take some events as logical parts of other events, giving logical overlap, not spatiotemporal overlap. But then the idea that these are different *events* (which are supposed to be particulars, or things that occupy regions of spacetime) rather than just different *properties* is still just as counterintuitive.

For Kim, when Suzy writes her paper with her left hand, multiple events occur: Suzy's writing of her paper, Suzy's writing of her paper with her left hand, etc. All of the occurrent events that 'include' Suzy's writing of her paper with her left hand (such as Suzy's writing her

paper with her left hand by writing on a keyboard, Suzy's writing her paper with her left hand by pecking with one finger on a keyboard, etc., if this was how Suzy actually wrote her paper) are effects of the accident. And all of these effects are among the causes of the publication of the paper, for they all 'include' Suzy's writing of the paper, and so they are sufficient under the laws to cause the publication of the paper.

David Lewis argues for a counterfactual approach combined with fine grained events, and denies that events must be fragile.⁸ For Lewis, actual events have both strong and weak essences, so that when an event that is C. Louise's sneezing loudly occurs, a second event (with a weaker essence) *also* occurs—C. Louise's sneezing. Events, as under Kim's view, are greatly multiplied.

For Lewis, the skiing accident is among the causes of the event of Suzy's writing her paper that is essentially a writing of her paper and is essentially with her left hand, as well as a cause of all the other events that occurred essentially with the left hand (i.e., Suzy's writing of the paper that is essentially a writing of the paper and essentially with the left hand, but only accidentally a writing using a keyboard, etc.) If the accident had not occurred, none of these events, as essentially specified, would have occurred. And (holding that the appropriate closest possible world in which we evaluate whether or not the publication of the paper occurs is a world in which no event that is essentially Suzy's writing of her paper occurs) each of the events that involve Suzy's writing of her paper that is essentially a writing of her paper, which includes many events caused by the accident, are causes of the publication of the paper. So Lewis must also accept the conclusion that the skiing accident is a cause of the publication.

The problems for the Kimian and Lewisian views highlight a side effect of their

multiplication of entities: many unexpected instances of what we might call a kind of redundant causation. For Kim, each event that includes Suzy's writing of her paper causes the event that has as its constitutive property the publication of the paper. (Of course, there are many other publications of the paper caused at the same time as the one we are focusing upon, but we'll ignore those events for simplicity's sake.) For Lewis, each event that involves Suzy's writing of her paper essentially causes a version of a publication of her paper that is essentially a publication. Multiplying events in either the Kimian or the Lewisian way gives counterintuitive results in examples like that of the accident, results in a new and implausible kind of redundant causation, and, as I shall argue below, is not necessary. The point is not that the size of the ontology of events is unacceptable, but rather that fine grained theories of events cannot be motivated by the claim that they improve the results given by their companion theories of causation.⁹

Coarse grained events. Donald Davidson holds that the correct theory of causation involves a regularity approach, and that events are individuated by the regions of spacetime they occupy.¹⁰ This view has the happy consequence that C. Louise's sneezing is the same event as the event of her sneezing loudly, and the event of her sneezing in the kitchen, etc. Such a view takes events to be coarse grained.

Defenders of coarse grained events seem to have a natural solution to cases like the skiing accident causing the publication. They can hold that Suzy's writing the paper with her left hand and Suzy's writing the paper are one and the same event, and use this claim to break the link between the skiing accident and the writing of the paper. A natural way to justify the claim

would be to point out that that the actual writing of the paper and the actual writing of the paper with the left hand occupy the same spatiotemporal region.

Those who hold a version of the regularity analysis could argue that the skiing accident is not a part of a sufficient condition for the writing of the paper, so is not a cause of it.

Counterfactual theorists could adopt a moderately robust view of events together with the coarse grained approach and argue that the writing of the paper would have been the very same event even if it had been written with the right hand, so if the skiing accident had not occurred, the writing of the paper would still have occurred. With either approach, since Suzy's skiing accident was not a cause of the existence of the paper, it was not a cause of her publishing the paper.

Something about this seems right—it seems to capture part of what we mean when we deny that the accident was a cause of the publication of the paper. But it can't be quite right, because the response can be used to generate problems. For suppose there was more to tell about the case of the skiing accident. Suzy's left hand, unused to writing, begins to cramp severely after finishing the paper. She visits a doctor, and spends large amounts of money on prescription drugs in order to dull the pain. It seems right to say that the writing of the paper with Suzy's left hand is among the causes of the pain in her left hand, and her visit to the doctor and expenditure of large amounts of money on prescription drugs is caused by the pain. Moreover, it seems right to say that the skiing accident to Suzy's right hand is among the causes of her writing the paper with her left hand and so a cause of the pain in her left hand.

But if we are coarse grained theorists, we have already decided in order to resolve earlier problems that the event of Suzy's writing the paper is the same event as her writing the paper with her left hand, and that the skiing accident did not cause the event of the writing of the paper. But if

the skiing accident did not cause the event of the writing of the paper when we determine what caused the publication, then neither did it cause the event of the writing of the paper when we determine what caused the costly prescription. The link in the causal chain between the skiing accident and the writing of the paper has already been—by stipulation—broken. But we *do* want to say that the skiing accident was a cause of the cramping hand and the costly prescription, and moreover, we want to say this because we think that the accident, by causing the writing of the paper to occur in a particular way—it was a writing with the *left* hand—was a cause of the subsequent cramping of the left hand (and thus of the costly prescription).

So we shouldn't try to preserve causal intuitions by breaking the link between the skiing accident and the writing of the paper. Further, common sense would have it that the writing of the paper (however it was written) is a cause of the publication of the paper, so it's not a good idea for the coarse grained theorist to try to deny the second link in the chain. I suspect that the problem is perfectly general: for any theory of events that individuates coarsely (i.e., for theories which do not individuate events as finely as can be with respect to properties), we can design puzzles like the one with the cramping hand.¹¹

This leaves advocates of event causation in a predicament, since whatever standard of individuation for events they impose, there exist clear problem cases. Distinguishing between events coarsely respects many of our causal judgments and commonsense views about individuation. But by holding that some changes (no matter how minor) in events do not result in a different event, we lose needed flexibility when just such a minor change affects the causal story that we want to tell. On the other hand, theories of events that are as fine grained as can be build every detail of the event into the causal story that we tell. Such precise specification gives us

many cases of spurious causation—unless we reject the transitivity of event causation—which is a steep price to pay for those of us who value our common-sense views about causation.

2. Property Instances

But there is a clear, intuitive solution to the case of the accident. To see it, set aside the idea that events are the only kinds of causes and effects, and think in terms of property instances with respect to events or individuals. With this in mind, consider our example. It seems right to say that because Suzy's accident involved Suzy's right hand being broken, it caused the writing of the paper to be a writing with the left hand. But the left handedness of the writing had nothing to do with the paper's being accepted for publication: as I've told the story, all that caused the paper to be accepted for publication was that the paper was written in the first place, that (presumably) it was a good paper, etc. So the accident's property of being an accident to Suzy's right hand did not cause the paper to have the property of being accepted for publication, but merely to have the property of being written with the left hand. This becomes clear when we think of causes and effects as instances of properties rather than as events. It's not that the case shows us causation isn't transitive, but rather that this is a case where the question of transitivity does not arise, since the property caused by the accident and the property that causes the publication are not the same property! Transitivity just doesn't apply.

The follow-up example, where the skiing accident is a cause of the cramping in my left hand, does involve transitivity: the skiing accident's property of being an accident to Suzy's right hand caused the writing of the paper to be a writing with the left hand, and the writing of the paper being a writing with the left hand caused the left hand to cramp. This solution conforms to

our intuitive understanding of what happened, and as a result generates the correct answer.

Examples like these give us a strong reason to accept the view that property instances rather than events *simpliciter* are the causal relata. We get further flexibility by allowing the property instances to involve individuals as well as events.

There is a related reason to take property instances as the causal relata. As the example of Suzy's accident helps to bring out, there is an essential tension between the goal of developing an adequate, acceptable theory of event individuation and the goal of developing an adequate theory of causation. It is not uncommon to adopt an extremely fine grained or fragile theory of events in order to get better results for one's theory of event causation (this is what usually motivates advocates of fine grained views), but the example of the accident shows us that such maneuvers bring trouble elsewhere.

By separating the two accounts, we can be free to develop our best theory of event individuation (and our best theory of individuation for individuals) apart from our theory of the causal relata and the causal relation. Even if the Suzy examples don't convince, the history of problems developing an acceptable theory of events as a companion to a theory of causation, together with the lack of parsimony (both of causal claims and events) of views like Kim's and Lewis's should be enough to motivate a switch to property instances.

There is one issue that needs immediate explication: how does my view that the causal relata are property instances differ from Kim's theory of events as property exemplifications? The obvious difference between the views is that my account shows us how to keep transitivity while explaining cases like Suzy's accident. But the root of the difference is less obvious: it is based on the fact that Kim's account requires that causation relate events, defined using constitutive triples,

not property exemplifications *simpliciter*.

To see how this is important, recall that Kim's theory of events is metaphysically substantial: the property instances referred to by the triples *really are* events. It is *not* the view that we should define the causal relata and then call whatever we end up with 'events'. Because it is a definition of events, not just any property that can be exemplified counts as an event: properties that modify other properties are not themselves events—presumably because modifiers are even less acceptable, intuitively, as real events than other properties are.

For Kim, our example of the property of being performed with the left hand would not count as an event itself but rather (as it should) as a modifier of the property of being a writing of a paper. It is hard to see a way to make sense of the idea that the property of being performed with the left hand could count as any sort of event. Kim handles modifiers by combining the modifier with an appropriate constitutive property to create a new constitutive property, which in turn defines a new event. So for Kim, the property of being performed with the left hand is not an event in its own right: the writing of the paper counts as an event, and the writing of the paper with the left hand is another event. In this way, Kim hopes to be able to make room for the ways in which we seem to want to take account of the importance of properties, including modifiers, in cases of causation, yet retain the idea that causation is a relation between events. Unfortunately, by creating new events for properties that are merely modifiers, Kim opens the door to counterexamples like that of Suzy's accident.

Kim could revise his view and claim that instead of creating a new event that is the writing of the paper with the left hand, we say instead that the event of the writing of the paper (defined by the triple of Suzy, the property of being a writing of a paper, and the time at which the writing

occurred) has the property of being performed with the left hand.¹² But while this seems sensible, it won't help Kim, since in order to get the right result in the case of Suzy's accident, it is the *property* of the event, being performed with the left hand, that has to count as an effect (and as a cause, with respect to the cramping). Since Kim holds that causation is a relation between *events*, the modifying property of being performed with the left hand can't be a cause or effect unless it somehow counts as an event in its own right.

By switching to property instances as causes and effects we sidestep the problems that Kim, Lewis, Davidson and others with related views face when developing theories about causation. Of course, we must still rely on an adequate theory of property individuation, but any of these theories of events and event causation must rely on a theory of properties as well. Removing a problematic level of analysis and solving difficult problems with transitivity should be sufficient motivation to make the switch.

I will use 'aspect' to refer to a property instance: an aspect is a particular's (a particular event or individual) having a property. Aspects are things that correspond one to one with thing-property pairs such that the property is had by the thing; so aspects are in an important sense part of the spatiotemporal world. Defined as such, aspects correspond to tropes (if there are any tropes), but the definition of 'property' is intended to be flexible: whether property instances involve exemplifications of universals, sets of particulars, states of affairs or tropes I need not say nor choose between.¹³ A few more details: aspects involving conjunctive properties can be causes or effects iff each conjunct is. The properties instanced must be suitably natural, so aspects that involve gruesome or disjunctive properties are not eligible to serve as causes or effects. The question of how extrinsic properties can be and still be paired with particulars to serve as causal

relata needs further investigation.¹⁴

Taking aspects as the causal relata does not exclude events as causes and effects, for it may be that when an aspect causes another aspect or group of aspects, the aspect or aspects that are caused are sufficient to imply or constitute an event. In some cases, if an aspect causes enough properties to be instantiated, or perhaps enough essential properties to be instantiated, we may say that the aspect causes an event in virtue of causing the particular to have those properties.¹⁵ If aspect c causes Billy to have the property of dying in a particular way at a particular time, this aspect might be sufficient for the instantiation of an event, viz. the death of Billy. If so, then aspect c is a cause of the event of the death of Billy, in virtue of being a cause of aspects sufficient to constitute the event. Depending on your standards for particulars, the aspects sufficient to constitute a particular may be many or few, common or rare. The power of the aspects account is that it will deliver the right causal judgment consistent with the standard of particulars adopted.

Counterfactual and regularity accounts based on aspects could rely on appropriate versions of the following definitions:

1. Counterfactual dependence of aspects: For any two distinct, actual events or individuals c and e , and logically distinct properties p and q : aspect e_q (e 's having q) is counterfactually dependent on aspect c_p (c 's having p) iff, had c occurred (existed) without p , then e would not have occurred (existed) with q . (I do not specify whether c without p is c or is numerically different from c .)¹⁶

2. Lawful entailment of aspects: For any two distinct, actual events or individuals c and e , and any

two logically distinct properties p and q , aspect c_p (c 's having p) lawfully entails aspect e_q (e 's having q) iff c 's exemplification of p is subsumed by the antecedent of the right law or laws which entail a consequent subsuming e 's exemplification of q . This is a souped-up kind of sufficiency, one that excludes properties had by particulars not linked to the effect via a law of nature from being counted as part of a lawfully sufficient condition for an effect.¹⁷

Under a regularity account where lawful entailment implies causation, being a breaking of the right hand is a cause of the left handedness of the writing. But it is the paper's property of being a good paper (rather than the property of being written with the left hand) that lawfully entails the paper's being published, so the accident's being a breaking of the right hand is not a cause of the publication.

Under a counterfactual analysis where dependence implies causation, the left handedness of the writing depends on the breaking of the right hand. However, the publication of the paper does not depend on the left handedness of the writing: it depends upon the property of being a good paper. Our transitivity puzzle is dissolved.

3. Reductive Analysis: Problems

Although the simple versions of the counterfactual and regularity accounts can be used to solve the transitivity puzzles we have been discussing, they cannot serve as theories of causation. Recent work on causation shows that problems with preemption, where potential causes c and b both occur, and each in the absence of the other can cause the effect e , yet it is intuitively clear that c is a cause of e and b is not, require a more sophisticated treatment of the causal relation.

Two kinds of preemption are responsible for the worst of the problems: *late preemption* and *trumping*. In cases of late preemption, the central issue involves the fact that the preempted cause *b* would have caused the very same effect *e*, but slightly later than the preempting cause *c* caused it. In these cases, the preempted causal chain is prevented by the occurrence of the effect itself, before the preempted cause can cause it.

C. Louise crouches, aiming for an unfortunate fly. Possum also crouches, aiming for the same fly. C. Louise pounces, and catches the fly. She then eats it. Possum, though agile, is heavier than C. Louise and so pounces more slowly, and the fly is eaten by the time he arrives. If C. Louise hadn't eaten the fly, Possum would have eaten it in the very same way, but just a few moments later. Counterfactual and regularity theories of event causation seem to have serious problems with examples like these: the (event that is the robust) catching of the fly does not depend counterfactually on C. Louise's crouching, and conversely, both C. Louise's crouching and Possum's crouching seem to be lawfully sufficient for the catching.

I have argued elsewhere that effects in cases of late preemption *do* depend on their causes, for when the effect occurs depends on whether the preempting cause occurred, but when the effect occurs does not depend on whether the preempted cause occurs.¹⁸ We can make a similar claim in terms of lawful entailment: the preempting cause, but not the preempted cause, lawfully entails the effect's occurring when it did.

Trumping examples give a new twist to the problem. In these examples, preempted cause *b* would have caused the very same effect *e*, but for the fact that a law specifies that if *c* occurs, *c* causes *e* and *b* does not. In these cases, *b* could have brought about the effect at the very same time, and with the very same properties, as *c* did.¹⁹

C. Louise crouches, aiming for another fly. Possum also crouches, aiming for the same fly. C. Louise jumps. Possum, who has been practicing, jumps a moment later, but his (newly acquired) agility makes him able to catch the fly at the same time as C. Louise. Unfortunately for Possum, there is a little-known law that states that flies, when pounced upon by multiple cats, are captured by the cat who jumps first. Since C. Louise jumps before Possum, she gets the fly. If C. Louise hadn't jumped, Possum would have captured the fly in the very same way and at the very same time. C. Louise's pounce, albeit through no intrinsic feline merit, trumps Possum's.

Trumping cases where the effect would have occurred at the very same time and in the very same way if it had been caused by the preempted cause are only of concern if we want our analysis to be able to handle action at a distance.²⁰ It seems to me that trumping is simply a new variant of early preemption, and cases not involving action at a distance can be solved using stepwise dependence.²¹ Nevertheless, the cases are quite interesting and I prefer an analysis that can handle them to one which cannot.

Unsurprisingly, counterfactual and regularity theories have problems with trumping: the catching of the fly does not depend on C. Louise's crouching, and conversely, C. Louise's crouching and Possum's crouching each lawfully entail the catching. Since there is no difference in when or whether the effect occurs if the preempting cause does not occur, the solution for late preemption cannot be straightforwardly applied to trumping cases.

Lewis sees that restricting dependence to when or whether an effect occurs if the cause occurs is too limiting²² and proposes that we define causation in terms of a pattern of dependencies between events. The relevant events are *alterations* of the events for which the causal relation is being evaluated. These alterations are either the actual *c* and *e* ('actual' refers to

the world of the example), or very fragile versions of, or alternatives to, c and e . The alterations are used to help us represent different ways c and e could have occurred (or different ways events that are very much like c and e could have occurred).

Under Lewis's account, an event c influences an event e iff there is a substantial range $c_1, c_2, c_3 \dots$ and $e_1, e_2, e_3 \dots$ of not-too-distant alterations of c and e such that if c_1 had occurred, then e_1 would have occurred, and if c_2 had occurred, e_2 would have occurred, etc. In this way we check to see if whether, when and how e occurs depends on whether, when and how c occurs. If there is a sufficiently large range of direct dependencies between alterations of e and c , then c influences e , so c causes e . To preserve transitivity, as in the original counterfactual analysis, Lewis takes the ancestral: c causes e iff there is a chain of stepwise influence from c to e .

The new analysis handles preemption problems elegantly: the preempting cause, if changed, would have caused changes in the effect. If C. Louise had batted the fly instead of catching it, or if she had caught the fly differently, the effect at the end of the corresponding causal chain would have been different: a batting instead of a catching, or a catching that occurred with more or less enthusiasm. Not so for Possum: whether he'd bat or how he'd catch would make no difference to the effect.

But the approach has two serious defects. First, the weakening of the dependence requirement to include more ways in which the effect can depend on the cause allows spurious causation. In the preemption cases, it seems right that if Possum had pounced earlier or (in the late preemption case) with more agility, the effect would have occurred earlier or in a different way. But if the effect depends on Possum's acts, then by the analysis above Possum's act counts as a cause. This is clearly an undesirable consequence.

Lewis recognizes this problem and attempts to minimize it by arguing that in most cases alterations in C. Louise's act make more of a difference than changes in Possum's act, and further that alterations in Possum's act, in the context of comparing this act to C. Louise's, are much more distant than alterations of C. Louise's act. Differences in degree or distance correspond to differences in influence and justify calling C. Louise's act, but not Possum's, the cause. Lewis argues further that if it turned out after taking degree and distance into account that there was not much difference between the influence of Possum's act and C. Louise's act, we would be justified in calling Possum a cause as well.

The view is defensible with respect to cases of trumping where both potential causes could bring about the very same effect. But in cases of late preemption we are not justified in calling Possum's preempted act a cause no matter how much influence he has. The fact remains that Possum's act does not lawfully entail the catching of the fly when it actually occurred, and intuitively lawful entailment is a necessary condition for causation. (Or at least it is for worlds like our own.) The problem with influence can be put more generally: consider two events *a* and *b* that we would normally take to be causally unrelated. Take *a* to be my body temperature and *b* to be the white pages of the manuscript strewn across my desk. If my body temperature were altered so that I radiated sufficient heat, the white paper would turn brown and curl at the edges. But surely the temperature of my body is not a cause of the whiteness of the paper: it does not lawfully entail the whiteness. Influence alone is not sufficient for causation.

The second problem with the account involves transitivity. To preserve transitivity (and to help solve some particularly worrying cases of early preemption), Lewis takes the ancestral of the influence relation to be causation. However, this move commits him to counterintuitive results with

respect to a flock of (supposed) counterexamples to transitivity.²³

Many of the cases have the same general form: some series of events (call this event pathway A), initiated by event *a*, starts to occur. If all the threatened events of A occurred, the series would culminate in the causing of event *c*. However, before the series of events that make up event pathway A have all occurred, event *b* causes event *c* via a different chain of events (call this event pathway B) connected to but different from A. For example, a train rushes towards Jesse James, who is tied to the tracks. If the train continues on its track, it will run over and kill James (event pathway A). A few minutes before the train runs over Jesse, his brother Frank flips a switch that causes the train to veer left onto a different track (event pathway B). Unluckily for Jesse, this track converges to the original track just before the spot where he is tied, and the train runs him over anyway. We assume that the train's diversion to the left hand track did not delay the train or change the event of Jesse's death in any way. (We might assume that the original track meandered a bit before converging with the left track, so that each track was exactly the same length.)

Now, intuitively, we want to deny that Frank's flipping of the switch was a cause of the train running over and killing James. But under the influence account—as under accounts of event causation generally—we cannot. The event of Frank flipping the switch influences the event of the train's being on the left hand track and the event of the train's being on the left hand track influences the convergence at the point just before it ran over Jesse, since there are alterations of the event of the train being on the left hand track just before it converges (namely, the alteration where the event is completely excised from history) that would result in the convergence not occurring. Since the event of the train's running on the track through the convergence point

towards Jesse influences his death, the event of the flipping of the switch is a cause of Jesse's death.

The case is related to our skiing accident case above, and Lewis's new account fails to handle it for the same reason that his earlier account fails it: a reliance on events as the causal relata allows too much information into the causal claim, and when this extra information is combined with transitivity, spurious causal results are easy to generate.²⁴

4. Aspect Causation

Both lawful entailment and influence go some way towards capturing the content of the causal relation. So why not combine the two? For it is because of a lack of lawful entailment that the influence account errs in counting certain events as causes, and it is because a influences effect e and b does not in cases where a trumps b that a regularity account errs in counting b as a cause. Each analysis alone is too permissive, but combined they can give us a simple, strong and elegant analysis of causation.

In the first part of this paper, I argued that property instances, not events, are causes and effects: problems with transitivity help to make this clear. Accordingly, I propose the following analysis of causation, based on the definitions (1) and (2) given in §2 and Lewis's definition of influence given in §3:

Aspect Causation: For any two aspects c_p and e_q :

- (i) if c_p lawfully entails e_q , and
- (ii) if c_p influences e_q , then

e_q is directly caused by c_p . Taking the ancestral of direct causation in order to give us causation, c_p is a cause of e_q iff e_q is directly caused by c_p or there is a chain of direct causation running from c_p to e_q .

The idea is this: take any aspect c_p that lawfully entails an effect e_q ; for each such aspect c_p , check to see that the effect exhibits dependence on the cause by checking for an appropriate pattern of dependence of aspects on aspects. Let alterations p_1, p_2, \dots , and q_1, q_2, \dots of property instances p and q be property instances which might be similar to but numerically different from p or q , and check to see if c_{p1} had occurred, then e_{q1} would have occurred, etc.²⁵ If and only if the appropriate pattern of dependence exists, c_p is a cause of e_q .

By including influence in my account, I allow for a certain amount of vagueness, but for far less than in the original influence theory. The account tightens up the influence theory in two major ways: it prevents anything from counting as a cause if it does not lawfully entail the effect, and it prevents illicit information from being included in causal claims by taking aspects rather than events (in the first instance) as the causal relata. It seems correct to say that both influence and lawful entailment capture part of the nature of causation, and my hope is that the two combined suffice for a simple and strong analysis of the (deterministic) causal relation in the actual world.²⁶

By requiring causes to lawfully entail their effects, we eliminate the major problems with spurious causation that Lewis faced. In our example of late preemption, C. Louise's having the property of pouncing in some particular way at some particular time counts as a cause of the fly's having the property of being caught in a particular way at a particular time. The properties of C. Louise's act, being a pouncing in such and such a way and at such and such a time lawfully entail

the fly's having the property of being caught when and how it actually was, and if the properties of C. Louise's act, being a crouching in such and such a way and at such and such a time were changed, the fly would not have been caught when and how it actually was. We can even say that since the properties of C. Louise's act caused the catching of the fly in a particular way at a particular time, the event of the fly's being caught was caused *simpliciter*.

But properties of Possum's act, being a pouncing with such and such a momentum and starting at such and such a time, do not lawfully entail the effect as it actually occurred. (Possum could not have pounced the way he did, when he did, and brought about the effect when it occurred the way it did.) Alterations of properties of his act that would affect the time of the fly-catching and thus give spurious dependence are not relevant to our evaluation of the situation, since these properties were what entailed his arriving too late in the first place. Likewise, my body's having a temperature of (about) 98.6 F does not lawfully entail the color of the paper on my desk, and so it cannot be counted as a cause no matter what the overall pattern of dependence looks like.

For those who resist analyses involving counterfactuals for fear of unwanted vagueness, my account should help to calm those worries as well. These worries can arise when determining the truth values of counterfactuals, for the selection function might specify more than one closest possible world, allowing conflicting truth values for the relevant counterfactuals. Unless we adopt a selection function that specifies only one closest world, ties in some cases are possible, and restricting causation to cases where there are no ties is *ad hoc*.²⁷ If we merely emend a counterfactual-based account of causation to say that a pattern of counterfactual dependence under some precisification of the similarity relation is sufficient for causation, then our

counterfactual constraint is too permissive, allowing causation whenever there is vagueness. On the other hand, requiring dependence under all precisifications is clearly too restrictive, preventing causation whenever there is vagueness.

But if we require lawful entailment for causation, then we can rely on the more permissive version of counterfactual dependence, yet eliminate most of the vagueness worries. Although we might have not have a pattern of dependence under all precisifications, intuitively, if we have dependence under one or more precisification of the similarity relation together with lawful entailment in the actual world we can say there is causation without being unduly permissive. There are additional benefits to combining lawful entailment with dependence, for lawful entailment brings in natural restrictions with respect to disjunctive, overly gruesome and logically necessary properties.²⁸ Including counterfactuals (and restrictions on backtracking counterfactuals) with an account of lawful sufficiency, on the other hand, helps us address problems for regularity accounts with distinguishing causes from effects and effects of a common cause.²⁹

The second major change is the use of aspects rather than events *simpliciter* as causes and effects. Return to our case of late preemption with C. Louise and Possum. It seems right to say that Possum's mass lawfully entails some of the properties of C. Louise's eating of the fly (those properties caused by the minute gravitational forces his mass exerts on C. Louise and the fly), and that if Possum's mass were changed these properties would also change. So Possum's mass counts as a cause of some of the properties of the eating. Is this a problem for our account?³⁰

I think rather that it is an advantage. In some contexts we might need to ask about what caused these minute gravitational effects, and then we would need to be able to cite Possum's

mass as among the causes. Note that, unlike Lewis, I need *not* hold that Possum's mass is a cause of the eating of the fly; rather, his mass is merely among the causes of the minute gravitational effects. Lewis simply denies that we should attend to these negligible differences: but in a context where we need to explain these differences such a denial is inexplicable. We need aspects to have an adequately precise account of the causal story.

Turning from the problems with simple spurious causation to those with transitivity, moving to aspects solves the so-called transitivity problem where the flipping of the switch has to count as a cause of Jesse's death. When we eliminate the extra, unwanted, information brought in by the events, we see that transitivity is not the culprit: events as the causal relata are responsible for the unwanted consequences.

As the story was told, when the switch is flipped, the only property (or properties) it causes are the train's being on the left-hand track, and any other properties to do solely with the train's being on the left (e.g., the train's having the property of passing the scenery on the left more closely than if it had traveled on the right). The flipping lawfully entails the train's traveling on the left, and the flipping of the switch influences the traveling on the left. The other properties of the train's travel (which are relevant to Jesse's death) while it travels on the left are caused by properties of the train's prior motion, mass, etc., not by the flipping of the switch: the flipping does not lawfully entail them, and they are not influenced by the flipping. (If the flipping had occurred differently or not at all, all of these properties would have occurred just as they did in the original story.)

Because the only (relevant) effect of the flipping is the train's traveling on the left, the flipping does not cause any of the properties of Jesse's death. The train's merely having the

property of traveling on the left does not lawfully entail the death of James. Nor does the property of traveling on the left influence properties of James's death, since even if the train had traveled, e.g., on the right, James's death would not have been affected.³¹

Note that, under the aspects account, if we were to tell the train story differently, the account would reflect the right changes in the results. Telling the story *very* differently might muddle our intuitions enough to make it unclear whether the flipping should count as a cause of properties of the death. In such a case, an answer either way, or no clear answer at all, is acceptable. Telling the story *slightly* differently, so that the train is somehow slowed or changed by taking the left track and the death happens later or is more violent than if the train had taken the right track, makes the flipping of the switch count among the causes of those properties (e.g., the time or manner of death)—just as it should.³²

Conclusion

By recognizing that property instances are the causal relata, we can conform to common sense; we can develop, unimpeded, the account of the individuation of events (and individuals) that best conforms to our antecedent beliefs; and we can allow our causal relata to be at least as fine grained as facts.

By combining this view with lawful entailment and influence, we arrive at a well developed theory of property causation (perhaps most useful in the context of current discussions of mental causation), we solve many of the most pressing problems with extant analyses of causation, and we have an account of causation that fits well with actual scientific practice. By this last claim, I mean that when doing science, we identify the potential causes eligible to bring about an effect

(those that lawfully entail the effect), and then alter the potential causes and observe the effect for changes. When doing metaphysics rather than empirical science, we evaluate close possible worlds to see if alterations of the potential cause (one which lawfully entails the effect) result in alterations of the effect. If we have the right sort of dependence, we have causation. If not, not. The account I propose is modeled after the ways we rely on to learn about causation in experimental contexts, so it is no surprise that it does such a good job of handling cases where we have clear causal intuitions and a full description of the situation. We thus clear away a number of problems that have hampered the development of the analysis of the causal relation.

¹I received helpful comments from many, but I am especially indebted to Ned Hall, Chris Hitchcock, David Lewis, Jonathan Schaffer, Stephen Yablo and an audience at the 1998 Australasian Association for Philosophy.

²The example is modeled on a nice example used in Michael McDermott's "Redundant Causation", *British Journal for the Philosophy of Science* 40 (1995), 523-44. McDermott uses his example to create trouble for a particular kind of counterfactual account that takes events to be coarse grained but does not discuss its broader implications.

³David Lewis, 'Causation as Influence', this issue.

⁴See David Lewis, 'Causation', *Philosophical Papers*, vol. 2 (Oxford: New York, 1986), 159-72, and 'Postscript to Causation', same volume, 172-213.

⁵Jaegwon Kim, 'Causation, Nomic Subsumption, and the Concept of Event', *Journal of Philosophy* 70 (1973): 217-36 and 'Events as Property Exemplifications', in M. Brand and D. Walton (eds.), *Action Theory* (Reidel: Dordrecht, Holland, 1980): 159-77, reprinted in Kim's *Supervenience and Mind: Selected Philosophical Essays* (Cambridge, 1993): 33-52.

⁶Kim, 'Events as Property Exemplifications': 45-6 (page numbers refer to *Supervenience and Mind*).

⁷Jonathan Bennett points this out as well in *Events and Their Names* (Hackett: Indianapolis, 1998), p. 83.

⁸Lewis, 'Events', *Philosophical Papers*, vol. 2 (Oxford: New York, 1986), 241-69.

⁹Stephen Yablo's work goes a great way towards handling many of the problems I discuss in this paper, but also relies on extremely strong essentialist claims and the ontology these claims imply. I prefer a more parsimonious account (if one can be had). Yablo develops his account in 'Mental Causation', *The Philosophical Review*, 101 (1992), 245-80, 'Cause and Essence', *Synthese*, 93 (1992), 403-49, and 'Seven Habits of Highly Effective Thinkers', *Proceedings of the 20th World Congress of Philosophy*, forthcoming.

¹⁰Davidson originally tried to individuate events by their causes and effects, but such a strategy requires a prior analysis of causation. Donald Davidson, 'Events as Particulars', 181-7, and 'The Individuation of Events', 163-87, in his *Essays on Actions and Events*, (Clarendon Press: Oxford, 1980). Quine argued that the strategy didn't succeed, and Davidson agreed. See 'Events and Reification', W.V. Quine, and 'Reply to Quine on Events', Donald Davidson, both in *Actions and Events: Perspectives on the Philosophy of Donald*

Davidson, eds. Ernest LePore and Brian McLaughlin (Blackwell: Oxford, 1985), 162-76. After writing this paper, I learned that Douglas Ehring, *Causation and Persistence*, (Oxford: New York, 1997), and Daniel Hausman *Causal Asymmetries* (Cambridge, 1998) raise similar transitivity worries for Davidson.

¹¹ There is one more move that the coarse grained theorist might try: following Davidson's 'Causal Relations', *The Journal of Philosophy*, 64 (1967), 691-703, he might say that 'The accident caused the publication of the paper' is true, but not causally explanatory. In other words, the claim would be that, strictly speaking, the accident caused the publication even though our singular causal statement doesn't mention any causally relevant properties of the accident as part of the description it gives of the cause (the accident). This seems as though it would allow the coarse grained theorist to accept the first link of the causal chain (as well as subsequent links). But this response is inadequate, for the skeptic can return with the point that no matter *how* the accident is described, the event of the accident has *no* properties that are causally relevant to the effect (the publication). Davidson's argument requires that the event that counts as the cause have *some* properties under *some* description that are causally relevant to the effect, even if these are unobvious (e.g., such as the properties of Suzy's birth that are causally relevant to Suzy's writing her paper).

¹² Kim, 'Events as Property Exemplifications': 44-5.

¹³ Donald C. Williams, 'On the Elements of Being', *The Review of Metaphysics* vol. VII (1953): 3-192, and Keith Campbell, *Abstract Particulars* (Basil Blackwell: Oxford, 1990), suggest that tropes should be the causal relata. Douglas Ehring, *op.cit.*, argues for 'persisting tropes' as causal relata, and Hausman, *op.cit.* argues that tropes aid causal explanation.

¹⁴ For the purpose of using aspects as part of a theory of causation, I leave the controversial question of how we define particulars aside by stipulating that the aspects don't have to be part of the structure we take to *be* the particular. Just as C. Louise's sneezing isn't part of the structure that is C. Louise (and unlike the way C. Louise's left front paw is part of the structure that is C. Louise), the aspects that are the causal relata need not be the structure we take to define them.

¹⁵ D. H. Mellor has a very nice discussion of (what he calls) the difference between causing and affecting particulars. See his *The Facts of Causation* (Routledge: London, 1995), chap. 12.

¹⁶ In my account of dependence, I am assuming a standard ordering relation on closeness of worlds based on Lewis's similarity ordering, in which propositions about property instances are eligible to be included amongst the relevant conditions for the evaluation of counterfactuals.

¹⁷ viz. Kim's account of subsumption in 'Causation, Nomic Subsumption, and the Concept of Event'. My account requires an adequate specification of the causal laws and the properties they refer to, but does not require allegiance to a particular theory of laws: one could rely upon the Dretske-Armstrong-Tooley view of nomic necessity, some sort of primitiveness claim, or the Mill-Ramsey-Lewis view that the right regularities are those which belong to the set which does the best job of systematizing and organizing physical information. The definition might change slightly depending on the theory of properties adopted: for simplicity, I've written the definition here for aspects taken as tropes.

¹⁸ L.A. Paul, 'Keeping Track of the Time: Emending the Counterfactual Analysis of Causation', *Analysis* 58:3 (1998), 191-8.

¹⁹ Jonathan Schaffer, 'Trumping Preemption', this issue.

²⁰ I do not claim that there can be no such thing as causation by action at a distance. Rather, once we eliminate the need for chains of events in our cases of causation, our intuitions change enough to motivate a change in the conditions necessary for causation: causation by action at a distance is different enough from ordinary causation to need an analysis of its own. (So cases which combine "ordinary" causation with causation by action at a distance will require a combined analysis, and it's no surprise if cases involving action at a distance are not handled by an analysis for contiguous causation.) I suspect that the analysis of causation by action at a distance will be a version of the lawful entailment analysis for aspects. Ned Hall, 'Two Concepts of Causation', *Philosophical Review*, forthcoming, argues that causation by action at a distance is a kind of causation that differs from the 'ordinary' kind, and develops an analysis based on minimally sufficient sets. Similarly, there are problems for causation by omission.

²¹ See Lewis's 'Causation' and 'Postscript to Causation' discussions of early and late preemption, and his stepwise solution to early preemption. In these papers Lewis lumps cases involving action at a distance in with late preemption (where the effect is delayed): I think this is a miscategorization.

²² Lewis, 'Causation as Influence'.

²³ 'Causation as Influence'. Lewis refers to these cases as 'Black-Red examples'. Also see Hall, 'Causation and the Price of Transitivity', this issue, who has some particularly ingenious cases. There are more variations of these kinds of cases than I can deal with here, including well known probabilistic versions. The earliest discussions of the (deterministic) cases that I've found are in Peter van Inwagen, 'Ability and Responsibility', *The Philosophical Review* 87 (1987), 201-24, William L. Rowe, 'Causing and Being Responsible for What Is Inevitable', and John Martin Fischer and Mark Ravizza, 'Responsibility for Consequences'. The latter two papers appear in Fischer and Ravizza, eds., *Perspectives on Moral Responsibility* (Cornell: Ithaca 1993), 310-21 and 322-47 respectively.

²⁴ Lewis's influence account also fails to handle the skiing example: the skiing accident influences the writing of the paper and the writing of the paper influences the publication.

²⁵ Following Lewis's account of alterations of events, I take no rigid stand on much we can alter a property for our account of influence: it is a matter of degree and context. I also leave vague how we decide which alterations are closer to the original p and which are more remote.

²⁶ I hope, with Lewis, that a suitably modified version will be able to handle indeterministic causation.

²⁷ Adopting Robert Stalnaker's account would allow us to select only one closest world. This does not solve our problem but merely shifts the worry, since according to depending on the selection function we adopt we could get different values for the counterfactual. In other words, truth or falsity of the counterfactual would be relative to the selection function. On Stalnaker's account we can superevaluate, so a counterfactual conditional is true (false) iff it is true (false) for every selection function, but the conditional is undefined in the case of vagueness. Robert Stalnaker, 'A Defense of Conditional Excluded Middle' in W.L. Harper, R. Stalnaker, and G. Pierce (eds.), *Ifs* (Dordrecht, 1980), 87-104.

²⁸ For those who find influence unacceptable or who are not concerned about trumping cases, I advocate a view that combines simple counterfactual dependence and lawful entailment (versions of (1) and (2) in §2) together with aspects.

²⁹ Lewis, 'Causation' shows how counterfactual based accounts can address these problems.

³⁰ Lewis discusses this problem in 'Causation as Influence', as does Jonathan Schaffer in 'Causes as Probability-Raisers of Processes, forthcoming.

³¹ Here I am relying on *traveling on the right* as the relevant property alteration to evaluate influence.

³² My solution in this section helps to resolve problems with explaining why Frank is not morally responsible for the death of James because he flipped the switch: merely flipping the switch didn't make him causally responsible for the death, so neither did it make him morally responsible.