The Metaphysics of Causation

First published Sun Feb 2, 2003; substantive revision Tue Jul 5, 2016

What must a world be like, to host causal relations? When the cue ball knocks the nine ball into the corner pocket, in virtue of what is this a case of causation?

Questions about the metaphysics of causation may be usefully divided into questions about the causal relata, and questions about the causal relation. Questions about the causal relata include the questions of (1.1) whether they are in spacetime (*immanence*), (1.2) how fine-grained they are (*individuation*), and (1.3) how many there are (*adicity*). Questions about the causal relation include the questions of (2.1) how causally related and causally unrelated sequences differ (*connection*), (2.2) how sequences related as cause to effect differ from those related as effect to cause or as joint effects of a common cause (*direction*), and (2.3) how if at all sequences involving causes differ from those involving mere background conditions (*selection*).

Philosophers have, of course, disagreed over all of these questions. In what follows, I shall survey some of the main arguments in the literature.

- 1. Relata
    - 1.1 Immanence
    - 1.2 Individuation
    - 1.3 Adicity
    - 1.3.1 Contrastivity
    - 1.3.2 Secondary Relata
- 2. Relation
    - 2.1 Connection
    - 2.1.1 Probability, Process, and Problems

1. Relata

Question: What are the causal relata? When the cue ball knocks the nine ball into the corner pocket, what are the terms of this causal relation? An account of the causal relata should reveal what sort of thing they are, how many of them there are, and what job each does. In short, it should reveal their *category*, *number*, and *role*.

Options: The standard view of the causal relata is that they are of the category of *event*, and that their number is *two*, in the roles of *cause* and *effect*. So on the standard view, when the cue ball knocks the nine ball into the corner pocket, there is said to be an (actual) event e_1 of the cue ball striking the nine ball, and an (actual, distinct) event e_2 of the nine ball sinking into the corner pocket, such that e_1 is cause and e_2 effect. The standard view, in short, holds that the causal relata are a pair of events.

The standard view has, of course, been disputed on all counts. As to category, while the standard view casts the causal relata as events (Davidson 1980a and 1980b, Kim 1973, Lewis 1986b), one also finds considerable support for *facts* (Bennett 1988, Mellor 1995), and occasional support for such other entities as *features* (Dretske 1977), *tropes* (Keith Campbell 1990), *states of affairs* (Armstrong 1997), *situations* (Menzies 1989a), and *aspects* (Paul 2000). Allegiances are

further complicated by disagreements over what events, facts, and these other creatures are.

As to number and role, while the standard view numbers the causal relata at two (Davidson 1980b, Mackie 1965, Lewis 1986a), one finds some support for contrastive views featuring *three* and even *four* relata (Hitchcock 1996, Woodward 2003, Maslen 2004, Schaffer 2005, Menzies 2007, Northcott 2008, Weslake forthcoming) with the additional term(s) playing the roles of *causal alternative* and/or *effectual difference*. One also finds some support for additional relata of a different sort, including *descriptions* (Anscombe 1975, McDermott 1995), *models* (Menzies 2004, Halpern and Pearl 2005, Hitchcock 2007a), and/or *default states* (Menzies 2004, McGrath 2005, Hall 2007, Hitchcock 2007a, Halpern 2008). (In what follows I will reserve "the relata" for events or facts or whatnot, and use "secondary relata" when I wish to speak of the prospect of descriptions or models or defaults as additional relata. This is purely for expository convenience.)

Category: What is the category of the causal relata? What sort of thing are they? An account of the category of the relata should first determine which characteristics differentiate events from facts from the others, and then identify which characteristics the relata must have.

In practice, there are two main differentiating characteristics that one finds invoked. The first is *immanence*. Events and the others are generally regarded as immanent, while facts are generally regarded as transcendent. That is, the event of Brutus's stabbing Caesar is something concrete that occurs at a particular spatiotemporal location (the Senate on the Ides of March), while the fact that Brutus stabbed Caesar is something abstract and non-spatiotemporal. The question then arises as to whether causation requires immanent or transcendent relata.

The second main differentiator that one finds invoked is *individuation*. Events are sometimes (though not always) held to be relatively coarse-grained, while facts and the others are held to be relatively fine. For instance, the event of John's saying "hello" may be reckoned the same as the event of John's saying "hello" loudly, while the fact that John said "hello" is different from the fact that John said "hello" loudly. The question then arises as to whether causation requires coarse or fine relata.

Fortunately, questions of the true metaphysics of events, facts, and the other candidates may be postponed here, and the questions of the immanence and individuation of the causal relata may be addressed directly. Thus the issue of the category of the relata may, in practice, be usefully replaced by two questions: whether the relata are immanent (Section 1.1), and how finely they are individuated (Section 1.2).

Number and Role: What are the number and role of the causal relata? How many causal relata are there, and what kind of job do they do? An account of the number and role of the relata should first formulate general determinants of the adicity of relations, and then apply these determinants to causation.

The view that there are two relata is widely assumed but seldom defended. Three main alternatives have been explored involving contrastivity. The first of these alternatives, inspired by van Fraassen's (1980) work on contrastive explanation, takes causal relations to include an *effectual difference*. On this view causal relations have the form: c causes e rather than e^*. The second main alternative, based on Hitchcock's (1993, 1995a, 1996) work on probabilistic causation, takes causal relations to include a *causal alternative*. On this view causal relations have the form: c rather than c^* causes e. The third main alternative, defended by Schaffer (2005), includes both a causal alternative and an effectual difference and so numbers the causal relata at four, yielding the form: c rather than c^*

causes e rather than e^*. The question then arises whether contrasts (for cause and/or effect) help resolve problems and paradoxes (Section 1.3).

Three further main alternatives have been explored positing secondary relata. The first of these, inspired by Anscombe's (1975) claim that causation is an intensional relation, takes causation to be relative to descriptions of the primary relata. On this view, causal relations have the form: c causes e relative to D, where D is an ordered pair of descriptions (for c and for e). The second of these alternatives, arising especially from Pearl's (2000) work on causal modeling, treats causation as relative to a certain sort of mathematical representation. On this view, causal relations have the form: c causes e relative to M, where M is an apt causal model of the situation. The third of these alternatives, which has roots in Hart and Honore's (1985) treatment of causation in the law, treats causation as relative to default states, which encode the states that are considered "normal" and "natural," as opposed to the deviant states. The simplest version of this view runs: c causes e relative to N, where N is an ordered pair of natural outcomes (concerning c and e). These views are all compatible. One could for instance hold that causal relations have the form: c causes e relative to D, M, and N. (Indeed one area of active research concerns the combination of causal modeling with defaults: see Blanchard and Schaffer (forthcoming) for a critical overview.) The question then arises as to whether any secondary relata are needed, or whether they all constitute an objectionable loss of objectivity, or an objectionable departure from the allegedly "obvious binarity" of causation (Section 1.3).

Presuppositions: Both the dispute over the category and over the number and role of the causal relata involve presuppositions of *uniqueness*. As to category, the dispute presupposes that there is a unique category of entity from which all causal relata must be drawn. Yet, it might be argued, ordinary language allows for the relata to be described in eventive

(imperfect nominal), factive (perfect nominal), and other forms (Mackie 1974, Vendler 1984, Bennett 1988). Why not take ordinary language at its word, and let a thousand relata bloom?

As to number and role, the dispute presupposes that there is a unique number that is *the* adicity. Yet, it might be argued again, ordinary language allows for causal attributions with and without causal alternatives or effectual differences (Hitchcock 1996). Why not take ordinary language at its word, and let causation go multigrade?

There are two main arguments in defense of uniqueness, the first of which is that it staves off *ambiguity* (Menzies 1989a). If there were four choices for two relata, it might seem that there would be 2^4=16 "causal" relations (and more if there were more choices and/or relata and/or adicities). That said, it is unclear why there couldn't be a single causal relation (univocally denoted by "causation") which allowed different types of relata. The identity relation, for instance, can relate items of any ontological category.

The second argument for uniqueness is that it precludes a mysterious *harmony* (Mellor 1995). If there were a plurality of event causes and fact causes and the like, some metaphysical harmony would be needed amongst them, for surely they could not conflict. That is, it seems that the event of the cue ball's striking the nine ball, and the fact that the cue ball struck the nine ball, must have comparable effects. But without a unique underlying causal relation, there would seem to be nothing keeping these effects aligned. That said, perhaps a plurality of causal relations could be harmonized, provided either (i) one were fundamental and the others derivative, or (ii) all were derivative from a common non-causal basis, such as the regularities among the events.

1.1 Immanence

Question: Are the causal relata immanent, or transcendent? That is, are they concrete and located in spacetime, or abstract and non-spatiotemporal?

This question is connected to the question of category. If the relata are transcendent, then they are facts. If they are immanent, then they are events, or one of the other candidates such as features, tropes, or situations.

In practice, one finds two main arguments on the question of immanence. First, there is the argument from pushing, which maintains that the relata must be immanent so as to push things around. Second, there is the argument from absences, which maintains that the relata must be transcendent so that absences can figure in causal relations.

Pushing: The main argument for immanence is that only immanent entities can *interact*. This argument is nicely summarized by one of its opponents, Bennett: "Some people have objected that facts are not the sort of item that can cause anything. A fact is a true proposition (they say); it is not something *in* the world but is rather something *about* the world, which makes it categorically wrong for the role of a puller and shover and twister and bender." (1988, p. 22; see also Hausman 1998) According to the pushing argument, only concrete spatiotemporal entities can be causes and effects.

There are two main responses to the pushing argument, the first of which is to find substitute immanent entities. These substitute immanents serve as pushers, and relate to the causal facts, while still being distinct from them. Bennett, in the immediate continuation of the above quote, recruits *objects* for just such a purpose: "That rests on the mistaken assumption that causal statements must report relations between shovers and forcers. I grant that facts cannot behave like elbows in the ribs, but we know what

items do play that role — namely, elbows. In our world the pushing and shoving and forcing are done by *things* — elementary particles and aggregates of them — and not by any relata of the causal relation." (1988, p. 22) Mellor (1995) offers a similar response, suggesting *facta* (the immanent truth-makers for facts) as the immanent basis for fact causation.

The second response to the pushing argument is to charge that it rests on a naive (pre-Humean) conception of causation as requiring some sort of metaphysical push or "oomph". If the causal relation is a mere matter of regularity, why can't the regularities hold between facts?

Absences: The main argument for transcendence is that *absences* can be involved in causal relations. Absences are said to be transcendent entities. They are *nothings*, non-occurrences, and hence are not in the world. Thus Mellor says, "For the 'C' and 'E' in a true causal 'E because C' need not assert the existence of particulars. They may deny it… They are negative existential statements, made true by the *non*-existence of such particulars, …" (1995, p. 132) Here Mellor is arguing that, in the case where rock-climbing Don does not die because he does not fall, Don's non-falling and non-dying are causally related, without there being any events or other immanent entities to relate.

There are two main responses to the absence argument, the first of which is to deny that absences can be causal. In this vein, Armstrong claims: "Omissions and so forth are not part of the real driving force in nature. Every causal situation develops as it does as a result of the presence of positive factors alone." (1999, p. 177; see also Beebee 2004a and Moore 2009) The theorist who denies absence causation may add some conciliatory codicil to the effect that absences stand in cause-like relations. Thus Dowe (2000, 2001) develops an account of ersatz causation (causation*) to explain away our intuitions that absences can be genuinely causal.

The second response to the absence argument is to deny that absences are transcendent. One way to do this would be to accept the existence of negative properties, and think of absences as events in which an object instantiates a negative property. Thus Don's instantiating non-falling at t_0 might be counted an immanent event, and a cause of the further immanent event of his instantiating non-dying at t_1. A second way to deny that absences are transcendent would be to take absence claims as merely a way to describe occurrences, as Hart and Honore recommend: "The corrective here is to realize that negative statements like 'he did not pull the signal' are ways of describing the world, just as affirmative statements are, but they describe it by *contrast* not by *comparison* as affirmative statements do." (1985, p. 38) Thus Don's not falling at t_0 may be identified with his clinging to the rock at t_0, and Don's not dying at t_1 may be identified with his surviving at t_1, which events are indeed causally related.

1.2 Individuation

Question: How are the causal relata individuated? That is, if r_1 and r_2 are causal relata, what are the conditions that determine whether $r_1 = r_2$?

This question is related to the question of category. If the relata are coarse-grained, then they are events of a certain stripe. If they are fine-grained, then they are facts, or one of the other candidates such as features, tropes, or situations, or else events of another stripe.

Combining the considerations of immanence and individuation, one might, as a first approximation, distinguish a square of possible views underlying the dispute over category:

|	Coarse-grained	Fine-grained
Immanent	Davidson	Kim, Lewis, Dretske, etc.
Transcendent	[unoccupied]	Bennett, Mellor

In the upper-right quadrant one finds, alongside fine-grained events theorists such as Kim and Lewis, virtually all those who have opted for a third way, including Dretske, Campbell, Armstrong, Menzies, and Paul. Indeed, virtually all the theorists who have rejected both events and facts have done so because they think that the relata must be immanent and thus not facts, but fine-grained and thus not events.

But the square is only a very rough first approximation, because granularity really comes in degrees. So to a second approximation, one might render granularity on a continuum, from the coarseness of Davidson's (1980c) view on which the causal relata are individuated by their causes and effects, to the moderate fineness of Kim's (1976) view on which the causal relata are individuated by their associated <object, property, time> triples, to the extreme fineness of Bennett's (1988) view on which the causal relata are individuated as finely as propositions.

The Davidson-Kim-Bennett continuum is still only an approximation though, because there really are more occupied points along a wider continuum. How fine-grained Kim's events actually are depends on how finely properties are individuated. If properties are taken in an abundant sense (individuated as finely as predicates, or at least as finely as necessarily coextensive predicates) then Kimian events are relatively fine-grained, whereas if properties are taken in a sparse sense (individuated by "the joints of nature") then the grain is coarser. How fine-grained Bennett's facts are depends on how finely propositions are individuated. If propositions are taken as Frege conceives them, then Bennett's facts are exceptionally fine-grained; whereas if propositions are taken as Russell conceives them, then Bennett's facts are comparable in granularity to

Kim's events with abundant properties. Moreover, Davidson's view is outflanked in coarseness by Quine's (1985; accepted by Davidson 1985) view that the causal relata are individuated by spacetime region. While even the Fregean variant of Bennett's view is outflanked in fineness by Dretske's (1977) view on which even focal differences (such as between *Mary's* kissing John, and Mary's kissing *John*) entail differences in relata. Putting this together yields the following picture:

Individuation: Coarse-Grained → Fine-Grained
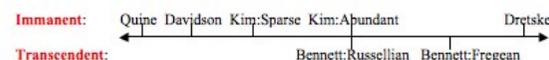
Further discriminations along the granularity dimension are, of course, possible.

In practice, one finds three main arguments on the question of individuation. First, there is the argument from causal differences, which maintains that the relata must be fine so as to mark differences in causal relatedness. Second, there is the argument from transitivity, which maintains that the relata must be fine so as to preserve transitivity. Third, there is the argument from methodology, which maintains that the relata must be coarse for general reasons of theoretical elegance.

Causal Differences: The first main argument for fineness is that fine differences can mark *causal differences*. To take an example discussed by Davidson, it makes sense to say, "The collapse was caused, not by the fact that the bolt gave way, but by the fact that it gave way so suddenly and unexpectedly" (1980b, p. 161; see also Kim 1976). This suggests that the bolt's snapping, and the bolt's snapping suddenly, must differ as causal relata. Or to take an example from Lewis (1986b), John's saying "hello" must differ from John's saying "hello" loudly, since only the former causes Fred to greet John in return, and only the latter is caused by John's state of

tension. According to the causal differences argument, the relata must be fine on pain of conflating conflicting causal relations.

There are three main responses to the causal differences argument, the first of which is that "cause" is *ambiguous* between causation and *explanation*. The causal differences argument is said to adduce explanatory differences only (Davidson 1980b and 1980d, Strawson 1985). Davidson integrates this response within a general account of causation and explanation, in which causation is an extensional relation that holds between coarse events, while explanation is an intensional relation that holds between coarse events under a description. So John's state of tension causes John's saying "hello," but the explanatory relation only holds when the "hello" is described in terms of its loudness. Our judgment that John's state of tension does not 'cause' John's saying "hello" is to be explained away as keyed into the explanatory idiom.

The second response to the causal differences argument is that "… causes…" is an *intensional* context (Anscombe 1975, Achinstein 1975 and 1983, McDermott 1995). Intensional contexts do not license substitution of co-referring terms *salva veritate*. Thus John's saying "hello" may refer to the same event as John's saying "hello" loudly, but substituting the one description for the other may still change the truth-value of the causal claim. If so then different *descriptions* of the same relata can induce causal differences.

The third response to the causal differences argument is that it *overextends*. It looks to require a fineness beyond what many of its proponents have envisioned. To take an example from Achinstein (1983), it may be true that "Socrates' *drinking hemlock* at dusk caused his death," but false that "Socrates' drinking hemlock *at dusk* caused his death." If so, the causal differences argument entails that Socrates's *drinking hemlock* at dusk, and Socrates's drinking hemlock *at dusk*, must differ as causal relata.

If so, then the causal differences argument ultimately requires the extreme fineness of the Dretskean view, which some may regard as a *reductio*.

Transitivity: The second main argument for fineness is that it preserves *transitivity*. To adapt an example from Woodward (1984; see also Ehring 1997, Paul 2000), suppose that Tom puts potassium salt in the fireplace (c), Dick then tosses a match in the fireplace, which results in a purple fire blazing in the fireplace (d), which then spreads and immolates Harry (e). The coarse-grained theorist looks to be committed to the following breakdown of transitivity: Tom's putting potassium salt in the fireplace causes there to be a purple fire blazing in the fireplace: c causes d; there being a purple fire blazing in the fireplace causes Harry's immolation: d causes e; but Tom's putting potassium salt in the fireplace does not cause Harry's immolation: c does not cause e. The fine-grained theorist may distinguish d_1: the fire becoming purple at region r, from d_2: the fire blazing at r. Now c causes d_1 (not d_2), d_2 (not d_1) causes e, and so the transitive inference to c's causing e would be blocked.

There are two main responses to the transitivity argument, the first of which is to bite the bullet. That is, one may accept that c does cause e: Tom's putting potassium salts in the fireplace does indeed cause Harry's immolation. Our intuitions to the contrary might be written off, as above, as confusing the causal and explanatory idioms. Or our intuitions to the contrary might be explained away on pragmatic grounds. As Lewis remarks in a related context, we are prone to "mix up questions of what is generally conducive to what with questions of what caused what," though, "Every historian knows that actions often have unintended and unwanted consequences." (2000, pp. 194–5)

The second reply to the transitivity argument is to deny that causation is transitive. There seem to be failures of transitivity that cannot be resolved by fine-graining. For instance, to borrow an example from Hall (2000, see

also Hitchcock 2001), suppose that the boulder begins to roll down the hill towards the hiker's head (*c*), which causes the hiker to duck (*d*), which in turn causes the hiker to survive (*e*). It seems that *c* causes *d* and that *d* causes *e*, yet it does not seem that *c* causes *e* or that slicing up *d* into different features or aspects or whatnot will help. If so, then transitivity is lost anyway.

Methodology: The main argument for coarseness is that it is *methodologically* preferable. Quine (1985, p. 167) charges fine-grained conceptions of the relata with invoking poorly individuated and unfamiliar entities, and recommends extremely coarse (spatiotemporal) individuation as principled and familiar in ontic commitments. And Davidson (1985) embraces Quine's view as both "neater" and "better" than even Davidson's own (1980b) previous view.

There are two main replies to the methodology argument, the first of which is that some of the fine-grained conceptions are perfectly principled. Kim's (1976) fine-grained conception of events as property instantiations, for instance, offers a precise criterion for individuation, namely:

$(\forall x)(\forall y)$ (If *x* and *y* are events, then *x* = *y* iff [Object(*x*) = Object(*y*) & Property(*x*) = Property(*y*) & Time(*x*) = Time(*y*)]).

(The generalization to *n*-ary relations is straightforward.) And the entities invoked (objects, properties, and times) should be perfectly familiar to all but the most austere nominalist. In fact, Quine himself admits not only that Kim's fine-grained conception is perfectly principled, but even that it "could still be accommodated in the ontology that I have accepted." (1985, p. 167)

The second reply to the methodology argument is that it is not clear that these indecencies in individuation or multiplication of entities should count for much. As to individuation, we accept physical objects without

clear individuation principles for those, so why hold events (or whatever the relata may be) to a stricter standard? As to multiplication, if one has a *reductive* fine-grained view (such as Lewis's 1986b view of events as transworld classes of regions), then there is no multiplication in one's basic ontology at all, since all the components already exist. And if one has a *nonreductive* fine-grained view, then the resulting multiplication may still be blameless. The real methodological sin is to multiply entities *without necessity*, so if there is a need for fine individuation (such as the causal differences and/or transitivity arguments might provide), then postulation of such entities is methodologically pure.

1.3 Adicity

Question: What is the number and role of the causal relata? That is, how many are there, and what kind of job does each do? This question may be usefully divided into the subquestions of whether to posit contrasts (Section 1.3.1), and whether to posit any secondary relata such as descriptions, or models, or defaults (Section 1.3.2).

1.3.1 Contrastivity

As to contrasts, one finds four main arguments. First, there is the argument from surface form, which maintains that the causal relata must be two to fit the surface form of utterances such as "the short circuit caused the fire." Second, there is the argument from determinacy, which maintains that the causal relata must be four for causal relations to be well defined. Third and fourth, there are the arguments from immanence-revisited and individuation-revisited, which maintain that the causal relata must be four to resolve the problems about absences and causal differences reviewed above (Sections 1.1 and 1.2, respectively).

All the arguments to be considered here call for either two or four relata (not three). Indeed, one might argue against contrastive three relata views that they preclude *causal chains*. In a causal chain the effect at the first link serves as the cause at the second. For this to be possible, cause and effect must be formally exchangeable: the same structure must flank both sides of the relation. Suppose the first domino knocks over the second, which then knocks over the third. The binary theorist can say that *c*: the toppling of the first domino, causes *d*: the toppling of the second; and that *d* in turn causes *e*: the toppling of the third domino. The quaternary theorist can say that *c* rather than *c**: the first domino's remaining upright, causes *d* rather than *d**: the second domino's remaining upright; and that *d* rather than *d** causes *e* rather than *e**: the third domino's remaining upright. But if there were contrasts on only one side of the relation, then no such chains could be constructed. The links would not match.

Surface Form: The main argument for binarity is that the *surface form* of causal claims reveals it. Causal claims like "the short circuit caused the fire" make no explicit reference to any contrasts. Such claims can be felicitously uttered out of the blue (in discourse initial position), and so do not require any antecedent contrast setting or presupposition fixing, either. This is presumably the root of the idea that causation is "obviously binary". Indeed this sort of consideration is most prominent in Davidson (1980b), who seeks the logical form of such surface-binary utterances. Relatedly, Davidson rejects the nearby notion of *causal relevance* because, "There is no room for a concept of 'cause as' which would make causality a relation among three or four entities rather than between two." (1993, p. 6)

There are three main replies to the surface form argument, the first of which is that contrastive surface forms exist too. For instance, one also finds such claims as "Pam's throwing the rock rather than the pebble caused the window to shatter," "Pam's throwing the rock caused the window to shatter rather than crack," and even the combined "Pam's throwing the rock rather than the pebble caused the window to shatter rather than crack." Thus surface form might seem equivocal (Hitchcock 1996).

The second main reply to the surface form argument is that surface form may be a reduced expression of a more complex logical form. For instance, "Ann prefers chocolate" may be used as a reduced expression of the proposition that Ann prefers chocolate *over vanilla*. Here the contrast does not need to be explicitly articulated, or even noted earlier in the conversation. One can make surface-binary preference claims out of the blue. So just as preference claims might seem to have a contrastive logical form beneath their binary surface, so causal claims too might be regarded in this light. At any rate, as the history of semantics has shown, it would be foolishly naive to think that ordinary language displays its logical form on its surface. In this vein Schaffer (2012) diagnoses a range of context-dependencies in causal discourse as being due to implicit contrastivity.

The third main reply is to go *revisionary*. Even if the logical form of causal attributions ultimately proves to be binary, logical form should not have the last word in metaphysics anyway, since it might predicate a property that we have theoretical reason to reject. For instance, "The rock is moving" might seem to predicate the property of absolute motion, but physicists have discovered that there is no such thing. So just as motion claims only make metaphysical sense when relativized to an inertial frame, one might think that causal claims only make metaphysical sense when relativized to contrasts.

Determinacy: The first main argument for quaternicity is that binary causal relations are *ill defined*. Suppose that Jane smokes moderately, and develops lung cancer. Does Jane's moderate smoking cause her lung cancer? Hitchcock says that there is no determinate answer unless one

fixes the causal alternative: "The solution to this puzzle is to deny that there is any such thing as *the* causal relevance of moderate smoking for lung cancer… Relative to heavy smoking, it is a negative cause of (prevents) lung cancer; relative to abstaining, moderate smoking is a positive cause of (causes) lung cancer… Relations of positive or negative causal relevance only hold relative to specific alternatives." (1996, p. 402) A parallel case could be made for needing to fix the effectual difference. Suppose that Pablo is choosing between blue, red, and green paint for his canvas. Does Pablo's choosing blue paint rather than red cause the canvas to be blue? Here one might say that there is still no determinate answer. Pablo's choosing blue paint rather than red causes the canvas to be blue rather than red, but does not cause the canvas to be blue rather than green. Thus it may be concluded that contrasts are required for both cause and effect, in order for causal claims to have determinate truth-values.

The main reply to the determinacy argument is that binary causal relations are well defined, after all. This reply should take the form of applying a binary account of causation to problem cases such as the smoking and painting cases, and simply reading off a truth-value, whatever it may be. For instance, one might think that a straightforward counterfactual account of causation, on which we check whether *e* would still have occurred had *c* not occurred, simply rules that Jane's smoking causes her lung cancer, and that Pablo's choosing blue paint causes the canvas to be blue, full stop.

Immanence Revisited: The second main argument for quaternicity *revisits immanence*, and maintains that additional argument places reconcile immanence with absence causation (Schaffer 2005). The reconciliation is attempted through treating absence-claims as setting the contrast to the associated occurrence. For instance, "the gardener's failing to water the flowers caused the flowers' wilting" is to be interpreted as: what the gardener actually did (viz., the actual event of his watching television) rather than watering the flowers (the non-actual event that is

the associated occurrence) caused the flowers to wilt rather than blossom. And this claim may well be true. In this way, all four of the relata may be treated as immanent entities, and absence causal claims may still come out true. Indeed, in this way absence causation requires no special provisions at all (which, as Dowe (2000) explains, is not the case on virtually any binary theory).

The main reply to the immanence-revisited argument is that immanence needs no revisiting. No reconciliation of absences and immanence is needed, and so no additional argument places are needed. This reply may take the form of denying immanence (that is, embracing facts), denying that absences are causal, or maintaining that there are immanent absences (Section 1.1).

Individuation Revisited: The third main argument for quaternicity *revisits individuation*, and maintains that additional argument places tame the causal differences argument (Schaffer 2005). The concern is that the causal differences argument overextends, in requiring that Socrates' *drinking hemlock* at dusk, and Socrates' drinking hemlock *at dusk*, differ as causal relata. The taming is attempted through treating focal differences as contrastive differences. Thus "Socrates' *drinking hemlock* at dusk" is to be interpreted as *c*: Socrates' drinking Hemlock at dusk, rather than *c**: Socrates' drinking wine at dusk (or some other contextually salient alternative to drinking hemlock); while "Socrates' drinking hemlock *at dusk*" is to be interpreted as *c*: Socrates' drinking Hemlock at dusk, rather than *c**: Socrates' drinking hemlock at dawn (or some other contextually salient alternative to occurring at dusk). And these different contrasts may induce different effects. In this way, focal differences may be allowed to yield causal differences, without having any implications for individuation, much less the extreme fineness of the Dretskean view.

The main reply to the individuation-revisited argument is that individuation needs no revisiting. No taming of the causal differences argument is required, and so no additional argument places are needed. This reply may take the form of maintaining the Davidsonian distinction between causation and causal explanation, of allowing that causation is an intensional relation, or of simply accepting the Dretskean view of the relata (Section 1.2).

1.3.2 Secondary Relata

As to secondary relata, one finds three main arguments. First, there is the argument from objectivity, which maintains that any relativization of causal relations to descriptions, models, or defaults is incompatible with the objectivity of causation. Second, there is the argument from the success of causal modeling techniques, which maintains that causation must be relative to a model so that modeling techniques may be exploited. Third, it has been argued that default relativity provides the best explanation for our intuitions across a range of problem cases. (The argument from surface form (Section 1.3.1) may also serve as a fourth argument against any secondary relata.)

Objectivity: Leaving aside linguistic issues of surface form, the main argument against descriptions, models, and/or defaults as secondary relata is that including these would compromise the objectivity of causation. When the cue ball knocks the nine ball into the corner pocket, it might seem that there is objective causation in the world. How we choose to describe the events, model the situation, or label certain outcomes as "default" or "deviant" seems besides the point. For instance, if a model M is used in which this interaction comes out as non-causal, one might well infer, *not* that the interaction is not causal relative to M, but rather that M itself is a poor model of the real, model-independent causal facts.

There are two main (related) replies to the objectivity argument, the first of which is to bite the bullet and deny that causation is objective. Such a response might be partly encouraged by the thought that causation seems to drop out of fundamental physics (2.1.2). The second related reply is to maintain that there is still objectivity in a wide range of surrounding notions. Thus Hitchcock invokes the notion of *token causal structure* for what causal models represent, and suggests that "we can afford to let judgments of token causation be infected by pragmatic criteria without giving up on the objectivity of causation generally: objectivity can be retained at the level of token causal structure" (2007a, p. 504).

Causal Modeling: Perhaps the main argument for model relativity begins with the success of causal modeling techniques, as developed in Spirtes, Glymour, and Scheines (1993), Pearl (2000), and Halpern and Pearl (2005), *inter alia*. These techniques have provided working algorithms for causal discovery, and yielded elegant accounts of type-level causal relations. It it only prudent, the argument then begins, to seek an account of (actual, token) causation using such techniques. Indeed it might well seem retrograde to do otherwise. But given that different models yield different causal verdicts, and given that there is no unique notion of a *canonical model* for a given situation (at least none yet developed), it might seem that the only remaining option is to relativize causal relations to models (Halpern and Pearl 2005, p. 85).

There are at least three replies to the causal modeling argument to consider. First, one might evince skepticism concerning causal modeling. Causal modeling certainly provides an excellent account of causal epistemology, but it is not obvious that one should draw any metaphysical conclusions from that. Second, one might attempt to define the notion of a canonical model. Third, one might provide an account of causation involving existential quantification over models (Hitchcock 2001) or even

universal quantification, so as to help oneself to modeling techniques without paying any further price in objectivity.

Problem Cases: Perhaps the main argument for description and default relativity, which also features as an argument for model relativity, is the utility of such additional relata in resolving problem cases. This style of argument covers a wide variety of claims. Just to mention a few examples, McDermott (1995) builds description-relativity into an overall account of causation that enjoys a fair amount of success. McGrath (2005) has suggested that our intuitions about which absences are causes may turn on expectations as to which are normal. Hall (2007) has shown that certain causally different cases may take isomorphic causal models, so that modeling techniques may themselves require something like a default/deviant distinction to discern these cases. And Hitchcock (2007a) uses models with assigned default states to define up the notion of a *self-contained causal network*, which he then puts to work in addressing many of the most difficult cases in the literature. (See Halpern and Hitchcock 2010 and forthcoming, as well as Blanchard and Schaffer forthcoming, for further debate about the need for a default/deviant distinction within causal modeling.) It is difficult to say much more about these arguments in general, without delving into all the details. Obviously further consideration of these arguments depends on detailed considerations as to how the theories in question fare versus their rivals, and also on delicate balancing of whether the solution preferred is worth the price of the additional relativity.

2. Relation

Question: What is the causal relation? When the cue ball knocks the nine ball into the corner pocket, what is the basis for this causal link? An adequate account of the causal relation should reveal where the causal lines run, which direction the causal arrows point, and what if anything distinguishes causes from mere background conditions. In short, it should reveal the basis for *connection*, *direction*, and *selection*.

Network Model: The causal relation is typically understood with reference to what Steward (1997; see also Beebee 2004) calls "the network model." The network model has two main components. First, it pictures the causal relation as a directed segment, and the causal relata as nodes. Second, the network model pictures history as a vast causal network. On the network model, given some realistic assumptions, the causal history of *e* forms a vast inverted tree (though one which ultimately narrows back down to the big bang):
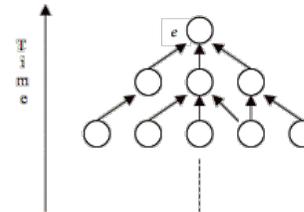
An account of connection is an account of the segments; an account of direction is an account of the arrowheads. Selection on the network model would consist in highlighting certain nodes:
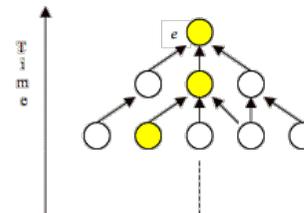
The network model without selection is implemented, for instance, in the neuron diagrams popularized by Lewis. In neuron diagrams, circles

doubly represent neurons that fire and events that occur, and arrows doubly represent synapses that stimulate and causation that obtains.

The network model without selection is also implemented in the directed acyclic graphs used in causal modeling to partially visualize the models (Spirtes, Glymour, and Scheines 1993; Pearl 2000). In directed acyclic graphs, nodes doubly represent variables with a range of values and occurrences with a range of alternatives, and links doubly represent functional determination via the structural equations and causal parenthood. (For further explanation of causal modeling, see the entry: causation: probabilistic.)

(Neuron diagrams and directed acyclic graphs are different implementations of the network model, though. Neurons represent events that occur or not: they are two-valued. Nodes represent variables that can take two, three, or even continuously many values. Synapses represent actual causation. Links represent lines of possible causal influence, and do not entail any actual causation between the actual values of their variables. See Hitchcock 2007b for comparisons and reasons to favor the causal modeling approach.)

Connection: What is the metaphysical basis for causal connection? That is, what is the difference between causally related and causally unrelated sequences?

The question of connection occupies the bulk of the vast literature on causation. One finds analyses of causation in terms of *nomological subsumption* (Davidson 1980d, Kim 1973, Horwich 1987, Armstrong 1999), *statistical correlation* (Good 1961 and 1962, Suppes 1970, Spirtes, Glymour, and Scheines 1993, Kvart 1997 and 2004, Pearl 2000, Hitchcock 2001), *counterfactual dependence* (Lewis 1986a and 2000, Swain 1978, Menzies 1989b, McDermott 1995 and 2002, Ganeri, Noordhof, and

Ramachandran 1996, Yablo 2002, Sartorio 2005), *agential manipulability* (Collingwood 1940, Gasking 1955, von Wright 1975, Price and Menzies 1993, Woodward 2003), *contiguous change* (Ducasse 1926), *energy flow* (Fair 1979, Castaneda 1984), *physical processes* (Russell 1948, Salmon 1984 and 1998, Dowe 1992 and 2000), and *property transference* (Aronson 1971, Ehring 1997, Kistler 1998). One also finds views that are *hybrids* of some of the above (Fair 1979, Dowe 2000, Paul 2000, Schaffer 2001, Hall 2004, Beebee 2004b), along with *primitivism* (Anscombe 1975, Tooley 1987 and 2004, Carroll 1994, Menzies 1996) and even *eliminativism* (Russell 1992, Quine 1966).

Fortunately, the details of these many and various accounts may be postponed here, as they tend to be variations on two basic themes. In practice, the nomological, statistical, counterfactual, and agential accounts tend to converge in the indeterministic case. All understand connection in terms of *probability*: causing is making more likely. The change, energy, process, and transference accounts converge in treating connection in terms of *process*: causing is physical producing. Thus a large part of the controversy over connection may, in practice, be reduced to the question of whether connection is a matter of probability or process (Section 2.1).

Direction: What is the metaphysical basis for causal direction? That is, what is the difference between sequences related as cause to effect and those related as effect to cause or as effect to joint effect of a common cause?

The standard view on the direction of causation is that it reduces to the direction of time: causes occur prior to their effects (Hume 1975, Kant 1965, Beauchamp and Rosenberg 1981). The temporal view has fallen into disfavor of late, and a number of alternatives have been suggested, for instance, that the causal direction is the direction of *forking* (Reichenbach 1956, Horwich 1987, Papineau 1993, Dowe 2000), *overdetermination*

(Lewis 1979), *independence* (Hausman 1998), and *manipulation* (von Wright 1975, Price 1991 and 1996, Woodward 2003). On these alternative views, the coincidence of the causal and temporal orders is merely a contingent feature of the actual world, or at least a typical feature of our patch of it. So the question arises as to whether the causal direction is the temporal direction, or something else (Section 2.2).

Selection: What is the metaphysical basis for causal selection? That is, what is the difference between cause to effect sequences involving real causes, and those involving mere background conditions?

The standard view on selection is that there is *no objective basis* for selection: selection is interest-driven, pragmatic, and unsystematic (Mill 1846, Lewis 1986a, Mackie 1974). This no-objective-basis view is occasionally challenged, and alternatives have been suggested such as *sufficient* versus *necessary* (Ducasse 1926) and *abnormal* versus *normal* (Hart and Honore 1985). The question thus arises as to whether there is any objective difference between "triggering conditions" and "mere background conditions" (Section 2.3).

Presuppositions: The dispute over the causal connection involves a presupposition of *uniqueness*; there has to be a unique metaphysical relation of *causal connectedness* for there to be a question of that relation's nature. Yet, it might be argued, ordinary language allows for a range of irreducible causal notions, such as "scrape," "push," "wet," "carry," "eat," etc. (Anscombe 1975, p. 68) Why not take ordinary language at its word, and adopt a form of causal pluralism on which 'cause' is viewed as providing a merely nominal grouping of metaphysically diverse relations?

One might defend uniqueness by arguing that scraping, pushing, wetting and the others have a real (more than nominal) unity. They have common statistical, counterfactual, predictive, explanatory, and moral implications.

Or one might defend uniqueness as a theoretical ideal. Even if our actual concept of causation is ultimately what Skyrms (1984) calls an "amiable jumble" of principles (see also Sober 1985, and Hall 2004), this might be taken as a call for conceptual revision. Why wallow in the jumble, instead of tidying it?
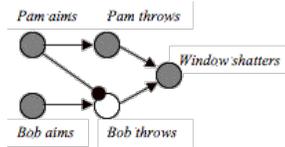
2.1 Connection

Question: What is the basis for causal connection? Is it a matter of probability, process, or some hybrid thereof? Is causal connection primitive and irreducible? Or is the entire notion merely a folk myth?

The probability and process views appear subject to a number of systematic problems, including the problems of *preemption* and *fizzling* for probability views, and *disconnection* and *misconnection* for process views (Section 2.1.1). These problems force the views to evolve and might be seen as motivating the hybrid, primitivist, and eliminativist alternatives (Section 2.1.2).

2.1.1 Probability, Process, and Problems

Preemption: One problem case for the probability view, on which the root idea of causation is that of making more likely, is the case of *preemption* (Good 1961 and 1962, Lewis 1986a, Menzies 1989b, Collins, Hall, and Paul 2004, Paul and Hall 2013 *inter alia*). Suppose that Pam and Bob each aim a brick at a window. Pam throws and shatters the window, while Bob holds his throw on seeing Pam in action. It seems that Pam's throw caused the window to shatter — her brick is what crashes through the glass. But it does not need to be the case that Pam's throw raised the probability of the shattering — if Bob is a more reliable vandal, then Pam's throw might even have made the shattering *less* likely. So it seems that probability-raising is not necessary for causation.

This sort of preemption (an *early cutting*) can be represented by the following neuron diagram:


A filled circle doubly represents a neuron firing and an event occurring. The line with a circular arrowhead doubly represents an inhibitory connection and a prevention.

Probability theorists have responded to preemption with three main strategies, the first of which is to *factor intermediaries* (Lewis 1986a, Menzies 1989, Eells 1991, Ramachandran 1997, Kvart 1997, Noordhof 1999, Pearl 2000, Yablo 2002). This approach looks to the course of intermediary events or probability evolutions to find some sense in which the preempting cause is indeed a probability-raiser. For instance, if one holds fixed the fact that Bob holds his throw, then it might seem that Pam's throwing becomes a probability-raiser after all. The second response that probability theorists offer to preemption is to *require precision* (Rosen 1978, Lewis 1986a, Paul 2000, Coady 2004). This approach looks to the exact time and manner of the events involved to try to find differences due to the preempting cause. For instance, if one considers the exact time and manner of the window shattering, then it might seem that Pam's throwing was a probability-raiser all along. The third response is to *consider intrinsic structure* (Lewis 1986a, Menzies 1996, Hall 2004) This approach looks to the intrinsic structure of the preempting process to try to find some sense in which the preempting process — even if it does not itself involve probability-raising — is at least intrinsically right for probability-raising.

It is unclear how far these three strategies extend. There is a vast literature on preemption involving a dizzying array of varied counterexamples and revised analyses. See Paul and Hall for an excellent guide to the "thick undergrowth of such examples" (2013, p. 5). It is fair to say that no currently available account is able to handle every case.

For a particularly difficult case, it has been argued that there can be preemption with no intermediaries to factor or differences to be precise about or intrinsic failure of the preempted backup, in cases of *trumping* (Schaffer 2000a, Lewis 2000). Suppose that the laws of magic say that the first spell cast on a given day matches the enchantment that midnight. Merlin casts a spell (the first that day) to turn the prince into a frog, Morgana casts a spell (the second that day) to turn the prince into a frog, and at midnight the prince turns into a frog. It seems that Merlin's spell caused the prince to turn into a frog — his spell was the first cast that day, and that's what the laws of magic identify as the relevant feature. Only Merlin's spell satisfies the antecedent conditions of the imagined law of nature. But it does not need to be the case that Merlin's spell raised the probability of the enfrogging — if Morgana is the more reliable wizard, then the chance of the enfrogging would have been greater has Merlin left the job to Morgana. It does not need to be the case that there are any intermediary events at all in the story — the magic might as well work directly. And it does not need to be the case that there would have been any differences in what befalls the prince had Merlin left it to Morgana. And finally, there does not seem to be any difference in "intrinsic aptness" between Merlin's and Morgana's spell. In fact the only relevant difference between these spells seems to be the extrinsic matter of which was cast first.

This sort of trumping preemption can be represented by the following neuron diagram (Paul and Hall 2013):
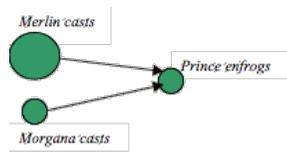
Here we suppose that neurons can fire in various colors (representing various possible spells), and that by law, when a neuron receives multiple stimulations, it fires in the color matching that of the biggest neuron (representing the first spell that day) stimulating it.

A third response that some probability theorists have advocated is to bite the bullet. Here it might be maintained that the effect occurred not because of, but *despite*, the preemptor (Eells 1991, Mellor 1995). Or it might be maintained that, at least in certain cases, the "preempted backup" is actually an overdetermining cause. For instance, Hitchcock (2011) argues that a contrastive approach to causation allows one to capture much of the intuitive asymmetry of trumping cases, while still counting the trumped event (Morgana's spell) as an overdetermining cause of the outcome (the prince becoming a frog).

Fizzling: A second problem case for the probability view is *fizzling*. Suppose that Pam and Fred each aim a brick at a window. Pam throws and shatters the window, while Fred simply walks away, or throws wide, or is preempted by Pam. It seems that Fred's aiming did not cause the window to shatter — Fred's brick never touched the glass. But it might be the case that Fred's aiming did raise the probability of the shattering — if there was some non-zero chance that Fred would succeed, and some non-one chance that Pam would succeed, then Fred's aiming might well have placed the window in greater danger. So it seems that probability-raising is not sufficient for causation.

The version of fizzling in which Fred is preempted can of course be represented by the early cutting neuron diagram above, with Fred standing in for Bob.

Probability theorists have responded to fizzling with the same strategies as for preemption, namely factoring intermediaries, requiring precision, or looking at intrinsic aptness (Menzies 1989b, Hitchcock 2004, Kvart 2004). For instance, if one holds fixed the fact that Fred throws wide, then it might seem that Fred's throwing is not a probability-raiser after all. Or if one considers the exact time and manner of the window shattering, then it might seem that Fred's throwing was not a probability-raiser after all. Or if one looks at the intrinsic character of the Fred process, then it might seem as if this process was intrinsically not apt to cause the window shattering.

It is unclear how far these three strategies extend. One might think that there can be fizzlings with no intermediaries to factor or differences to be precise about or intrinsic defectivness of the fizzled non-cause, in cases of *overlapping* (Schaffer 2000b). Suppose that Merlin casts a spell with a .5 chance of turning the king and prince into frogs; Morgana casts a spell with an independent .5 chance of turning the prince and queen into frogs; and the king and prince, but not the queen, then turn into frogs. It seems that Morgana's spell did not cause the prince to turn into a frog — the fact that the queen was unaffected shows that Morgana's spell fizzled. But it is the case that Morgana's spell raised the probability of the prince turning into a frog. Moreover it need not be the case that there are any intermediary events in the story at all (the magic might as well work directly). It need not be the case that there would have been any differences in what befalls the prince had Morgana's spell taken hold rather than Merlin's. And it need not be the case that Morgana's spell was intrinsically defective in any way (indeed it is only the external circumstance of the queen being unaffected that reveals that Morgana's spell fizzled).

This sort of overlapping case can be represented by the following neuron diagram:
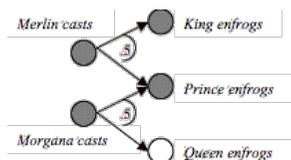
Here an arc with a number represents a conjunctive effect with a certain probability.

Preemption and fizzling cases together might be taken to show that probability-raising (however interpreted and refined) is systematically unable to provide necessary or sufficient conditions for causation. Moreover, these cases might suggest that connection is a matter of processes (be they physical or magical) rather than probabilities. The preempting cause and the effect are linked by a process, while the fizzled non-cause and the effect are not — just look at the diagrams. As Armstrong writes: "Where there is an arrow in a diagram showing that one neuron brings it about that another neuron fired, or is rendered incapable of firing, take it that here there is a genuine two-term relation of singular causation holding between cause and effect. Where there is no such arrow, deny that there is any such relation. This is the open door" (2004, p. 446).

But the process view faces problem cases of its own.

Disconnection: One problem case for the process view, on which the root idea of causation is that of physical connection, is *disconnection* (Ehring 1984, Schaffer 2000c, Lewis 2004, Hall 2004). Suppose that Pam catapults her brick through the window rather than throwing it. Then it seems that Pam's catapulting the brick causes the window to shatter — can it really matter here whether Pam catapults the brick or throws it? But

there need be no process connecting Pam's releasing the lever and the flight of the brick through the window — no relevant energy-momentum flow, track of mark transmission, or persisting trope connects them. Rather what is happening here is that the cocked catapult is prevented from launching by the catch, and Pam's releasing the lever prevents the catch from preventing the launch — the catapult is thus unleashed. The process of launch is purely internal to the catapult.

This sort of disconnection case can be represented by the following neuron diagram:
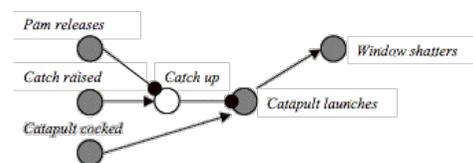
The main reply that is made to disconnections is to deny that they are genuinely causal. In this vein, Aronson says: "Consider a weight that is attached to a stretched spring. At a certain time, the catch that holds the spring taut is released, and the weight immediately begins to accelerate. One might be tempted to say that the release of the catch was the cause of the weight's acceleration. If so, then what did the release of the catch transfer to the weight? Nothing, of course." (1971, p. 425; see also Dowe 2001, Hall 2004) Indeed, since disconnections look to involve preventions of would-be-preventers, and since prevention looks to involve absences, one might reject causation by disconnection for the general reason that there is no absence causation (Moore 2009). For instance, one might deny that any causal connection can run through the absence of the catch.

Misconnection: A second problem case for the process view is *misconnection* (Hitchcock 1995b, Dowe 2000, Schaffer 2001). Suppose that Pam throws her brick through the window, while innocent Tom

watches in dismay, or sprays purple paint in the air through which Pam's brick passes. Then it seems that Tom's watching or paint-spraying does not cause the window shattering. But there is a process connecting Tom's watching or paint-spraying to the shattering. When Tom watches there will be photons connecting him to the shattering. When Tom sprays paint at the brick there will be a track of purple paint from Tom's spray can to the brick to the window. (Misconnections might be subdivided into *micro-connections*, which are of the wrong magnitude such as the photons, and *pseudo-connections*, which are of the wrong sort such as the paint.)

One finds two main replies to misconnections in the literature, the first of which is to bite the bullet. In the case of the photon connection from Tom's watching to the window shattering (and micro-connections generally), this might be regarded as causation of such negligible proportions that it is understandable that we might neglect it. Our intuitions to the contrary might also be written off, Davidson style, as confusing the causal and explanatory idioms.

The second main reply one finds to misconnections is to fine-grain the processes involved. In the case of the paint connection from Tom's spraying to the window shattering, the line of paint persistence and the line of brick flying through the window might be regarded as distinct and merely coincident (Dowe 2000). In this way it might be denied that there is a genuine process connecting the misconnecting non-cause and the effect.

Disconnection and misconnection cases together might be taken to show that process-linkage (however interpreted and refined) is systematically unable to provide necessary or sufficient conditions for causation. Moreover, these cases might suggest that connection is a matter of probabilities rather than processes. The disconnecting cause (such as Pam's release of the catapult) and the effect are linked by probability,

while the misconnecting non-cause (such as Tom's bystanding or spray painting) and the effect are not. So the probability theorist might claim revenge. But one might also draw the larger moral that the probability and process views are at best aspects of a bigger picture. This might inspire the search for a hybrid view.

2.1.2 Hybrids, Primitivism, and Eliminativism

Hybrids: Given the intuitive plausibility of the probability and process views, and the systematic problems each encounters, some recent theorists have sought reconciliation. Hybrid views aim to synthesize the probability and process views, capturing what is intuitively right about both ideas while resolving their problem cases.

The most obvious hybrid views simply conjoin or disjoin the probability and process views, or posit an ambiguity between the notions (Hall 2004). With respect to the above four problem cases, however, this might seem unpromising. Requiring both probability and process will resolve fizzlings and misconnections (Salmon 1997), as the former involves no process-linkage and the latter no probability-raising. But for the same reason it will not resolve preemptions and disconnections, as the former involves no probability-raising and the latter no process-linkage. (The conjunctive theorist might still avail herself of one of the above replies to preemption and disconnection.) Likewise requiring either probability or process will resolve preemptions and disconnections, but not fizzlings and misconnections. (The disjunctive theorist might still avail herself of one of the above replies to those problems.)

More sophisticated hybrid views attempt to *integrate* the notions of probability and process, and not merely conjoin or disjoin them. Thus Fair (1979) ultimately moves from an energy flow view to a view that understands connection in terms of counterfactuals about energy flow. And

Schaffer (2001) suggests a generalization of this approach, on which causal connection is understood in terms of the probabilities of processes. This sort of hybrid view might seem to resolve all of the above problem cases. The preempting cause and disconnecting cause do raise the probability of the process that produces the effect; the fizzled non-cause and misconnecting non-cause do not.

However this sort of hybrid view has trouble with "hybridized" problem cases such as the following (Schaffer 2001). Suppose that Pam throws a brick through the wall of an aquarium, preempting Bob from doing the same. The aquarium shattering then causes the soaking of the carpet, by preventing the glass from preventing the water from spilling. This is a preemption case fed into a disconnection case. It seems that Pam's throw caused the soaking of the carpet — her brick is what broke the aquarium. But it need not be the case that Pam's throw raised the probability of the process producing the soaking — if Bob is a more reliable vandal, then Pam's throw might have even lowered the chance of the spillage process, by preempting Bob. And it might also be the case that Bob's aiming raised the probability of the process producing the soaking — Bob's aiming might have raised the chance of the spillage process, by threatening the aquarium.

(Such a case also makes trouble for the Hall style "two concepts" view, since Pam's throw turns out to satisfy *neither* of the two concepts of causation Hall reports.) This sort of preemptive disconnection case can be represented by the following neuron diagram:
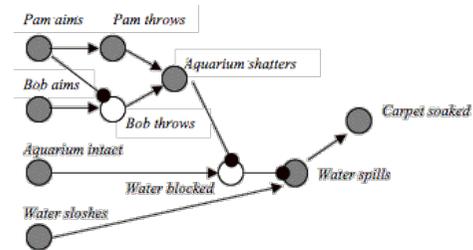
The current literature on causation is now suffused with complex hybrid cases, including cases of preemptive prevention (McDermott 1995, Collins 2000), disconnections inside larger chains (Hall 2004), and a variety of devious preemption variants (Paul and Hall 2013). No known account of causation — hybrid or otherwise — gets all of these cases right. So it is unclear what hybrid accounts gain in the end.

Primitivism: The problems encountered by attempts to analyze causal connection provide one main argument for primitivism. There seems to be a pattern of failure, which might suggest that causation is simply unanalyzable.

The second main argument for primitivism is that causation is *too central* to reduce. The probability and process accounts (and by extension hybrids) are said to be inevitably circular, because the notions of probability and process cannot be understood without reference to causation. As to probability, each of the nomological, statistical, counterfactual, and agential versions of the theory might be thought to harbor causal notions. The causal relation might be required to distinguish real laws from accidental generalizations (Armstrong 1997), to distinguish which background conditions must be held fixed in statistical assessment from which may vary (Cartwright 1983), to distinguish which background conditions may be held fixed in counterfactual supposition from which may vary (Kvart 1986), and to understand the notion of agential

intervention (Hausman 1998). As to process, it might be thought that a process is nothing more than a causal sequence; in the words of Sayre: "The causal process, continuous though it may be, is made up of individual events related to others in a causal nexus… It is by virtue of the relations among members of causal series that we are enabled to make the inferences by which causal processes are characterized." (1977, p. 206)

Indeed, the primitivist might add that the notion of event (or whatever the causal relata is taken to be) cannot be understood without reference to causation, because properties themselves are individuated by their causal roles (Shoemaker 1980 and 1998, Ellis 1999). As Carroll remarks in this regard, "With regard to our total conceptual apparatus, causation is at the center of the center." (1994, p. 118) So one might think that analysis is impossible because we have no more basic concepts.

As a third main argument for primitivism, it has been argued that there are worlds that differ purely causally (Armstrong 1983, Tooley 1987, and Carroll 1994). Suppose that it is a law of magic that all spells have a .5 chance of taking hold, that Merlin and Morgana each cast a spell to turn the prince into a frog at midnight, and that the prince turns into a frog at midnight. Then it may be intuited that there are three distinct possibilities: one in which only Merlin's spell caused the enfrogging, a second in which only Morgana's spell did it, and a third in which both did. These possibilities feature the same laws and pattern of events. So one might claim to intuit that causation is ontologically fundamental.

But there also three main arguments against primitivism, the first of which traces back to Hume and maintains that primitivism conflicts with the existence of causal knowledge. After all, one might argue (in a Humean vein) that all we can observe are sequences of events; as such, we could never come to know any facts about causal connection if connection is anything over and above such sequences. To this the primitivist might

reply, either that primitive connections can be directly observed, at least in certain favorable cases such as willing or pressure on the body (Anscombe 1975, Strawson 1985, Fales 1990, Armstrong 1997); or that primitive connections can be theoretically inferred via inference to the best explanation (Tooley 1987).

The second main argument against primitivism is that primitive causation is a spooky sort of primitive modality. To the extent that causal relations are supposed to be *necessary connections* in nature, and to the extent that one generally favors the reduction of the modal to the occurrent, one will have a general reason to resist any causal primitivism.

The third main argument against primitivism is that it leads to eliminativism. For if the options are irreducible causation or none, one should ask whether 'none' might be the better choice. For if science provides the criterion for which basic contingent entities one ought recognize, then the question must arise as to whether one could do science without any causal primitive whatsoever. One might conclude that our folk notion of causation must either reduce or face elimination.

Eliminativism: The final view of causal connection to consider is the eliminativist view, as trumpeted by Russell: "The law of causation,… is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm." (1992, p. 193) The eliminativist views the causal concept as a naively animistic projection of agency onto the world, to be superseded in a sophisticated scientific scheme.

The main argument for eliminativism is that science has no need of causation. The notion of causation is seen as a scientifically retrograde relic of Stone Age metaphysics. As Russell claims, "In the motions of mutually gravitating bodies, there is nothing that can be called a cause, and nothing that can be called an effect; there is merely a formula." (1992,

p. 202, see also Quine 1966) The differential equations of sophisticated physics are said to leave no room for causes, or at least to have no need of them.

Russell's argument might be effective against a primitivist treatment of causation, but the reductivist may reply that causation still reduces to scientifically respectable entities. In this respect, 'event,' 'law,' 'cause,' and 'explanation' are in the same boat. These nomic terms serve to allow a systematic understanding of science; they do not themselves appear in the equations. From this perspective, Russell's argument might seem akin to the foolish claim that math has eliminated the variable, because the term 'variable' does not appear in the equations!

The main objection to eliminativism is that causation is too central to eliminate. Causation, according to various contemporary philosophers, is required for the analysis of metaphysical concepts such as persistence, scientific concepts such as explanation and disposition, epistemic concepts such as perception and warrant, ethical concepts such as action and responsibility, mental concepts such as functional role and conceptual content, and linguistic concepts such as reference. Elimination is not just unjustified; it would be catastrophic. So it might seem that the arguments against primitivism and eliminativism bring one back to a reductive account of causation, and thus back again to probability and process, and their descendants.

Though perhaps there is some middle ground between the hopeless task of conceptual analysis on the one hand, and the conceptual disaster of primitivism/eliminativism on the other. For one should not confuse the prospects for conceptual analysis with the prospects for ontic reduction. It could well be that our concept of causation — something in the mind — is not definable via other concepts (in any way that would count as an informative conceptual analysis). Yet it could also be that the causal

relation — an entity out in the world — is not a fundamental constituent of reality. This middle position would claim to explain both the failures of conceptual analyses, and the disappearance of causation from fundamental physics. (See Schaffer 2007, pp. 872–873) for some further discussion.)

2.2 Direction

Question: What is the basis for causal direction? Is it a matter of temporal direction, or something else such as the direction of forking, the direction of overdetermination, the direction of independence, or the direction of manipulation? Is causal directedness a primitive, irreducible affair? Or is the belief in causal directedness merely a folk myth, or perhaps a projection of our experience as human agents onto an undirected world?

In practice, one finds six main arguments on the question of whether to identify causal direction with temporal direction. First, there is the argument from *bilking*, which maintains that the causal order must be the temporal order, or else the effect might occur but the cause then get prevented. Second, there is the argument from *time travel*, which maintains that the causal order must not be the temporal order because of the possibility of time travel. Third, there is the argument from *simultaneous causation*, which maintains that the causal order must not be the temporal order because of the possibility of cause and effect being contemporaneous. Fourth, there is the argument from *joint effects*, which maintains that it will not help to analyze the causal order as the temporal order because there are joint effect cases in which there is a temporally ordered connection without causation. Fifth, there is the argument from *physics*, which maintains that the causal order must not be the temporal order because of various physical hypotheses involving backwards causation. Sixth, there is the argument from *the causal theory of time*, which maintains that the causal order is the temporal order, but only

because the temporal order is to be analyzed in terms of the causal order, not *vica versa*.

Bilking: The main argument for the causal order being the temporal order is the bilking argument (Black 1956). The argument is explained by Mackie, with reference to a backward causal hypothesis that a drawing made by an alleged clairovoyant on Monday might be caused be a pattern made Tuesday: "But on every occasion, after the drawing is made, it is possible that someone or something should intervene so that the corresponding pattern fails to be produced. Consequently, it cannot on any occasion be the pattern that is responsible for the details of the drawing: the precognition hypothesis must be false even for those occasions when the device is not stopped, when the pattern is actually produced and turns out to be just like the drawing." (1974, p. 178) So temporally backwards causation is said to be impossible.

There are two main replies to the bilking argument. The first, due to Dummett (1964), is to note that the argument only applies to cases in which human intervention is possible. What is there to prevent backwards causation when human intervention is ruled out?

The second main reply to the bilking argument is that it involves an incoherent mix of determinism and indeterminism. If the world is deterministic, then the bilking intervention is impossible, as it will already be fixed that the later cause will occur. If the world is indeterministic, then the bilking intervention is possible but no longer problematic, as the case will then reduce to one in which the earlier event (e. g., the clairvoyant's drawing) is an uncaused indeterministic eruption.

Time Travel: The first main argument against the causal order being the temporal order is that temporally backwards causation is possible in cases such as *time travel*. It seems metaphysically possible that a time traveler

enters a time machine at time t_1, thus causing her to exit the time machine at some earlier time t_0. Indeed, this looks to be nomologically possible, since Gödel has proved that there are solutions to Einstein's field equations that permit looping pathways: "By making a round trip on a rocket ship it is possible [in worlds governed by Einstein's field equations] to travel into any region of the past, present, and future and back again, exactly as it is possible in other worlds to travel into distant parts of space." (1949, p. 560)

There are three main replies to the time travel argument. The first reply is that time travel is *incoherent*. A variety of incoherencies might be alleged here, including the incoherency of changing what is already fixed (causing the past), of being both able and unable to kill one's own ancestors, or of generating a causal loop and thus a reflexive relation of "self-causation", or of generating inconsistent probability assignments (Mellor 1995). The Gödelian proof might be dismissed as a mere mathematical artifact, not reflecting any possible situation.

The second reply to the time travel argument is that time travel may still take place through *locally* forward causal steps. Indeed, this is exactly what happens in the nomologically possible cases discovered by Gödel; spacetime is topologically structured in such a way that a series of locally forward steps produces a globally backwards path. This is compatible with the causal order being the temporal order, at least at each particular step. It may be that both the causal order and the temporal order can fail to possess global orientation.

The third reply to the time travel argument is that any alleged case of time travel is open to forward *redescription*. Instead of the time traveler entering the machine at t_1 and exiting at t_0, the same situation may be redescribed in terms of the spontaneous creation at t_0 of one individual,

and the spontaneous disappearance at t_1 of another, with merely coincidental correlations between their various mental and physical states.

Simultaneous Causation: The second main argument against the causal order being the temporal order is that *simultaneous causation* is possible. Indeed, it might seem that simultaneous causation occurs in the actual world, for instance when an iron ball depresses a cushion (Kant 1965, Taylor 1966, Brand 1980).

The main reply to the simultaneous causation argument is that the cases appearing to exemplify it are *misdescribed* (Mellor 1995). The iron ball takes time to depress the cushion, and in general all bodies take time to communicate their motions. There are no perfectly rigid bodies, at least in any nomologically possible world. Without the intuitive support of this sort of case, the simultaneous causation argument may be charged with begging the question. At this point methodological issues about the relevance of conceivable but physically impossible cases may arise.

Physics: The third main argument against the causal order being the temporal order is the argument from physics. Physicists of the past century have entertained a variety of theories positing backward causation, including the Wheeler-Feynman theory of radiation, Feynman's tachyon theory and his theory of positrons as electrons moving backwards in time, and de Beauregard's "quantum handshake" explanation of the violation of the Bell inequalities. While none of these theories enjoys much credence today, that they were serious physical hypotheses (at least once) seems to establish that they at least might have been true (Horwich 1987, Dowe 2000).

There are two main replies to the argument from physics, the first of which is to *dismiss* these theories. Perhaps these theories are all in the end false or even subtly incoherent. Here the possibility of forward redescription

alluded to with time travel is salient, in that the defender of the temporal order may hold that forward redescription is always possible, and always preferable.

The second reply to the argument from physics is that it *overextends*. There may be no coherent account of the causal order compatible with all these theories. In particular, the backwards causal model of the violation of the Bell inequalities postulates a backwards causal arrow that, it would seem, is neither the tine of a past-open fork, nor a special overdeterminer of the future, nor a handle to manipulate the past. So the argument from physics might culminate in a general *tu quoque*.

A different sort of argument from physics looks not to theories positing backwards causation, but rather to the lack of any relevant asymmetries (temporal or otherwise) in advanced physics. From this it is inferred that reality itself may be temporally symmetric, and that any sense of a *direction* to causal relations is due to a projection of our experience as agents (Price 1996 and 2007). Alternatively it might be argued that physics does supply a real direction — perhaps via Albert's "Past Hypothesis" (Albert 2000, Loewer 2007, Kutach 2007) — but one that is only contingently connected to the direction of time.

Joint Effects: The third main argument against the causal order being the temporal order is the problem of joint effects (Lewis 1986a). Suppose that the fall in atmospheric pressure at t_0 causes both the dip in the barometer at t_1 and the storm at t_2. Then the dip in the barometer and the storm are causally connected and temporally ordered, yet this is not causation, but rather a case of joint effects of a common cause.

This case may be represented by the following neuron diagram:
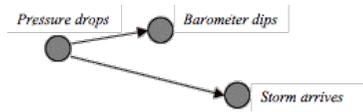
Here the left-to-right order represents the temporal order.

There are two main replies to the argument from joint effects, the first of which is to add some further test for a joint effect structure, such as the *screening off* test (Reichenbach 1956, Suppes 1970, *inter alia*). The causal direction is then taken as the unscreened temporal direction.

The second main reply to the argument from joint effects is to restrict the temporal order view to cases of direct connection. As the neuron diagram suggests, joint effects are only indirectly connected, via their common cause (Horwich 1987). If one can identify the direct connections and apply the temporal order to these only, one would have an arrow from the drop in atmospheric pressure to the dip in the barometer as well as to the storm, but no arrow from barometer to storm. This would match the diagram.

Causal Theory of Time: A final argument to consider is the argument that the temporal order is to be analyzed in terms of the causal order, not *vice versa* (Kant 1965, Reichenbach 1956, Mellor 1981). This argument cuts both ways, in the sense that it entails that the causal order is the temporal order (*contra* the arguments from time travel, simultaneous causation, and physics), but also entails that the causal order cannot be *based* on the temporal order on pain of circularity.

The main reply to the argument from the causal theory of time is, of course, to reject the causal theory of time. The temporal direction will have to be understood in other terms, perhaps in terms of intrinsic physical asymmetries such as involved with entropy or with the decay of the neutral kaon, or perhaps simply taken as primitive (Maudlin 2007). Though to the extent that one takes the direction of primitive, there is a

temptation to put this primitive to work to the extent possible, including understanding causal direction.

2.3 Selection

Question: What is the basis for causal selection? Is the distinction between real causes and mere background conditions merely an arbitrary and unsystematic affair? Or is there a metaphysical basis for selection having to do with sufficiency versus necessity, abnormality versus normality, or anything else?

(It should be noted that selection is widely associated with the idea of "the cause". This is perhaps a mistake. We often select multiple causes which act together. For instance, in a case of *joint causation* when the four movers collectively lug the piano up the stairs, it would be natural to select each individual mover's efforts as a real cause of the piano reaching the second floor (thereby selecting four real causes), all the while demoting various factors like the presence of the staircase to the status of being background conditions. 'The cause', relative to a given context, simply refers to something like the most salient cause at the context, just like 'the dog' refers to something like the most salient dog at the context; that is all a matter of what 'the' means and nothing at all to do with causation, or with dogs for that matter.)

In practice, one finds four main arguments on the question of selection. First, there is the argument from caprice, which maintains that our actual practice of selection is too capricious to have any real basis. Second, there is the argument from predictability, which maintains that our actual practice of selection is too predictable to be without a real basis. Third, there is the argument from inseparability, which maintains that we have no concept of causation that subtracts selection. Fourth, there is the argument

from adicity revisited, which alleges that additional relata might reconcile the preceding three arguments.

Caprice: The main argument for the no-basis view is Mill's argument from caprice: "Nothing can better show the absence of any scientific ground for the distinction between the cause of a phenomena and its conditions, than the capricious manner in which we select from among the conditions that which we choose to denominate the cause." (1846, p. 198) Mill's argument has won the field, and is echoed by contemporary authors such as Lewis: "We sometimes single out one among all the cause of some event and call it 'the' cause, as if there were no others. Or we single out a few as the 'causes', calling the rest mere 'causal factors' or 'causal conditions'… We may select the abnormal or extraordinary causes, or those under human control, or those we deem good or bad, or just those we want to talk about. I have nothing to say about these principles of invidious discrimination." (1986, p. 162) Thus selection is now generally dismissed as groundless, and theorists seek to isolate some pre-selected, egalitarian conception of causation.

Predictability: The main argument against the no-basis view maintains that our selections are too predictable to be without a basis. This point has been made by Hart and Honore, who write: "In most cases where a fire has broken out the lawyer, the historian, and the plain man would refuse to say that the cause of the fire was the presence of oxygen, though no fire would have occurred without it: they would reserve the title of cause for something of the order of a short-circuit, the dropping of a lighted cigarette, or lightning… In making this distinction it is plain that our choice, though responsive to the varying context of the particular occasions, is not arbitrary or haphazard." (1985, p. 11)

But what could this distinction between causes and conditions be? Ducasse maintains that it is between sufficient causes and necessary

conditions: "As a matter of established usage, 'cause' is contrasted with 'condition' in a serviceable and clearly stated manner: The *cause* of a phenomenon is a change in its antecedent circumstances which was *sufficient* to bring it about. A *condition* of a phenomenon, on the other hand, is a change, or more frequently a state, of its antecedent circumstances which was *necessary* to its having occurred when it did." (1969. p. 19) But it is difficult to see how this captures our selection of the short-circuit over the presence of oxygen, as each factor seems necessary and neither sufficient.

Hart and Honore maintain that abnormal situations and free actions are causes, while normal situations and non-agential factors are conditions: "In distinguishing between causes and conditions two contrasts are of prime importance. These are the contrasts between what is abnormal and what is normal in relation to any given thing or subject matter, and between a free deliberate action and all other conditions." (1985, p. 33) This seems to do better with respect to the short-circuit (abnormal) versus the presence of oxygen (normal), but at the price of such vagueness that one might think it only verbally distinct from the no-basis view.

Inseparability: A further argument against the no-basis view is that we have no concept of causation without selection. As Hart and Honore put it: "The contrast of cause with mere conditions is an inseparable feature of all causal thinking, and constitutes as much the meaning of causal expressions as the implicit reference to generalizations does." (1985, p. 12, see also Schaffer 2005). The upshot of this argument is that the no-basis view deprives us of any intuitive grasp on the notion of cause. For how are we to judge whether or not certain cases, such as the problem cases reviewed above or any others, involve causation or not, if our judgments are infected with a component of unsystematic caprice?

Lewis writes: "I am concerned with the prior question of what it is to be one of the causes (unselectively speaking). My analysis is meant to capture a broad and nondiscriminatory concept of causation." (1986a, p. 162) But it is not obvious that we *have* any such concept as Lewis seeks. Or at least, it is not obvious that our intuitions about causation can provide any *evidence* concerning this "broad and nondiscriminatory concept", if our intuitions are shot through with selection effects.

Adicity Revisited: A final argument to consider on the question of selection *revisits adicity* (Section 1.3), and maintains that additional causal relata might reconcile caprice and predictability, and explain inseparability (Schaffer 2005 and 2012). What is capricious about selection is that different speakers, in different conversational contexts, will disagree about what is a cause and what is a background condition. If one does not know what inquiry a speaker is pursuing, one may well find her selections capricious. What is predictable about selection is that, once conversational context is fixed, one can expect widespread agreement about causes versus conditions. If one knows what inquiry a speaker is pursuing, one will find her selections predictable. This might suggest that what is varying capriciously is which contrasts are in play in a given inquiry, and what is predictable is what counts as a real cause *relative to the contrasts in play*.

Mackie (a binary theorist) speaks of the *causal field* with reference to which causal selection is made: "A causal statement will be the answer to a causal question, and the question 'What caused this explosion?' can be expanded into 'What made the difference between those times, or those cases, within a certain range, in which no such explosion occurred, and this case in which an explosion did occur?' Both causes and effects are seen as differences within a field; anything that is part of the assumed (but commonly unstated) description of the field itself will, then, be automatically ruled out as a candidate for the role of cause." (1974, p. 35) The contrastivist may offer a natural implementation of Mackie's notion of

a causal field, understood now as the aspects of the situation that are assumed present and for which alternative (/contrasts) are not considered. This is particularly apt in understanding selection against the backdrop of causal models, in which only certain events are represented at all via variables (events that are represented via variables are outfitted with a range of alternative values, while events that are not represented at all are not).

Causal selection, on all views, is a reflection of which alternatives are in question. For the quaternary theorist, selection is determined by the values of the contrast relata (including the matter of whether any contrasts at all are being considered for a given factor). It is thus an inseparable component of our causal concept. For the person who thinks that causation is relative to a causal model, selection may be determined in a comparable way by the range of events modeled by variables (which are outfitted with a range of alternatives). On these views the very notion of a causal connection is only well defined in light of contrasts and/or models, and these additional relata explain selection.

Bibliography

Achinstein, Peter, 1975, "Causation, Transparency, and Emphasis," *Canadian Journal of Philosophy*, 5: 1–23.

——, 1983, *The Nature of Explanation*, Oxford: Oxford University Press.

Albert, David, 2000, *Time and Chance*, Cambridge, MA: Harvard University Press.

Anscombe, G. E. M., 1975, "Causality and Determination," in E. Sosa (ed.), *Causation and Conditionals*, Oxford: Oxford University Press, pp. 63–81.

Armstrong, D. M., 1997, *A World of States of Affairs*, Cambridge: Cambridge University Press.

——, 1999, "The Open Door," in H. Sankey (ed.), *Causation and Laws of*

Nature, Dordrecht: Kluwer Academic Publishers, pp. 175–85.

——, 2004, "Going through the Open Door Again: Counterfactual versus Singularist Theories of Causation," in J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*, Cambridge, MA: The MIT Press, pp. 445–57.

Aronson, Jerrold, 1971, "On the Grammar of 'Cause'," *Synthese*, 22: 414–30.

Beauchamp, Tom and Alex Rosenberg, 1981, *Hume and the Problem of Causation*, Oxford: Oxford University Press.

Beebee, Helen, 2004a, "Causing and Nothingness," in J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*, Cambridge, MA: The MIT Press, pp. 291–308.

——, 2004b, "Chance-changing Causal Processes," in P. Dowe and P. Noordhof (eds.), *Cause and Chance: Causation in an Indeterministic World*, London: Routledge, pp. 39–57.

Bennett, Jonathan, 1988, *Events and their Names*, Indianapolis: Hackett Publishers.

Black, Max, 1956, "Why Cannot an Effect Precede its Cause," *Analysis*, 16: 49–58.

Blanchard, Thomas and Jonathan Schaffer, forthcoming, "Cause without Default," in H. Beebee, C. Hitchcock, and H. Price (eds.), *Making a Difference*, Oxford: Oxford University Press.

Brand, Myles, 1980, "Simultaneous Causation," in P. van Inwagen (ed.), *Time and Cause: Essays Presented to Richard Taylor*, Dordrecht, D. Reidel Publishing, pp. 137–53.

Campbell, Keith, 1990, *Abstract Particulars*, Oxford: Basil Blackwell.

Carroll, John, 1994, *Laws of Nature*, Cambridge: Cambridge University Press.

Cartwright, Nancy, 1983, *How the Laws of Physics Lie*, Oxford: Clarendon Press.

Castañeda, Hector-Neri, 1984, "Causes, Causity, and Energy," in P.

French, T. Uehling, Jr., and H. Wettstein (eds.), *Midwest Studies in Philosophy* (Volume IX), Minneapolis: University of Minnesota Press, pp. 17–27.

Coady, David, 2004, "Preempting Preemption," in J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*, Cambridge, MA: The MIT Press, pp. 325–40.

Collingwood, R. G., 1940, *An Essay on Metaphysics*, Oxford: Clarendon Press.

Collins, John, 2000, "Preemptive Prevention," *Journal of Philosophy*, 97: 223–34.

Collins, John, Ned Hall, and L. A. Paul, 2004, "Counterfactuals and Causation: History, Problems, and Prospects," in J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*, Cambridge, MA: The MIT Press, pp. 1–58.

Davidson, Donald, 1980a, "Actions, Reasons, and Causes," orig. 1963, in *Essays on Actions and Events*, Oxford: Clarendon Press, pp. 3–19.

——, 1980b, "Causal Relations," orig. 1967, in *Essays on Actions and Events*, Oxford: Clarendon Press, pp. 149–62.

——, 1980c, "The Individuation of Events," orig. 1969, in *Essays on Actions and Events*, Oxford: Clarendon Press, pp. 163–80.

——, 1980d, "Mental Events," orig. 1970, in *Essays on Actions and Events*, Oxford: Clarendon Press pp. 207–27.

——, 1985, "Reply to Quine on Events," in E. LePore and B. McLaughlin (eds.), *Actions and Events: Essays on the Philosophy of Donald Davidson*, Oxford: Basil Blackwell, pp. 172–6.

——, 1993, "Thinking Causes," in J. Heil and A. Mele (eds.), *Mental Causation*, Oxford: Clarendon Press, pp. 3–17.

Dowe, Phil, 1992, "Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory," *Philosophy of Science*, 59: 195–216.

——, 2000, *Physical Causation*, Cambridge: Cambridge University Press.

——, 2001, "A Counterfactual Theory of Prevention and 'Causation' by

Omission," *Australasian Journal of Philosophy*, 79: 216–26.

Dretske, Fred, 1977, "Referring to Events," in P. French, T. Uehling, Jr., and H. Wettstein (eds.), *Midwest Studies in Philosophy II*, Minneapolis: University of Minnesota Press, pp. 90–9.

Ducasse, C. J., 1926, "On the Nature and Observability of the Causal Relation," *Journal of Philosophy*, 23: 57–68.

Dummett, Michael, 1964, "Bringing About the Past," *Philosophical Review*, 73: 338–59.

Eells, Ellery, 1991, *Probabilistic Causality*, Cambridge: Cambridge University Press.

Ehring, Douglas, 1986, "The Transference Theory of Causality," *Synthese*, 67: 249–58.

——, 1997, *Causation and Persistence*, Oxford: Oxford University Press.

Ellis, Brian, 1999, "Causal Powers and Laws of Nature," in H. Sankey (ed.), *Causation and Laws of Nature*, Dordrecht: Kluwer Academic Publishers, pp. 19–34.

Fair, David, 1979, "Causation and the Flow of Energy," *Erkenntnis*, 14: 219–50.

Fales, Evan, 1990, *Causation and Universals*, London: Routledge Press.

Ganeri, Jonardon, Paul Noordhof, and Murali Ramachandran, 1996, "Counterfactuals and Preemptive Causation," *Analysis*, 56: 219–25.

Gasking, Douglas, 1955, "Causation and Recipes," *Mind*, 64: 479–86.

Gödel, Kurt, 1949, "A Remark about the Relationship between Relativity Theory and Idealistic Philosophy," in P. Schilpp (ed.), *Albert Einstein: Philosopher-Scientist*, La Salle: Open Court, pp. 557–62.

Good, I. J., 1961, "A Causal Calculus I," *British Journal for the Philosophy of Science*, 11: 305–18.

——, 1962, "A Causal Calculus II," *British Journal for the Philosophy of Science*, 12: 43–51.

Hall, Ned, 2000, "Causation and the Price of Transitivity," *Journal of Philosophy*, 97: 198–222.

——, 2004, "Two Concepts of Causation" in J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*, Cambridge, MA: The MIT Press, pp. 181–204.

——, 2007, "Structural Equations and Causation," *Philosophical Studies*, 132: 109–136.

Halpern, Joseph, 2008, "Defaults and Normality in Causal Structures" in G. Brewka and J. Lang (eds.), *Principles of Knowledge Representation and Reasoning*, Palo Alto: AAAI Press, pp. 198–208.

Halpern, Joseph and Judea Pearl, 2005, "Causes and Explanations: A Structural-Model Approach— Part 1: Causes," *British Journal for the Philosophy of Science*, 56: 843–87.

Hart, H. L. A. and A. M. Honore, 1985, *Causation in the Law*, orig. 1959. Oxford: Clarendon Press.

Hausman, Daniel, 1998, *Causal Asymmetries*, Cambridge: Cambridge University Press.

Hitchcock, Christopher, 1993, "A Generalized Probabilistic Theory of Causal Relevance," *Synthese*, 97: 335–64.

——, 1995a, "The Mishap at Reichenbach Fall: Singular vs. General Causation," *Philosophical Studies*, 78: 257–91.

——, 1995b, "Salmon on Explanatory Relevance," *Philosophy of Science*, 62: 304–320.

——, 1996, "The Role of Contrast in Causal and Explanatory Claims," *Synthese*, 107: 395–419.

——, 2001, "The Intransitivity of Causation Revealed in Equations and Graphs," *Journal of Philosophy*, 98: 273–99.

——, 2004, "Do All and Only Causes Raise the Probabilities of Effects?" in J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*, Cambridge, MA: The MIT Press, pp. 403–18.

——, 2007a, "Prevention, Preemption, and the Principle of Sufficient Reason," *Philosophical Review*, 116: 495–532.

——, 2007b, "What's Wrong with Neuron Diagrams?" in J. K. Campbell,

M. O'Rourke, and H. S. Silverstein (eds.), *Causation and Explanation*, Cambridge, MA: The MIT Press, pp. 69–92.

——, 2011, "Trumping and Contrastive Causation," *Synthese*, 181: 227–40.

Horwich, Paul, 1987, *Asymmetries in Time*, Cambridge, MA: The MIT Press.

Hume, David, 1975, *An Enquiry Concerning Human Understanding*, orig. 1748. Oxford: Clarendon Press.

Kant, Immanuel, 1965, *Critique of Pure Reason*, orig. 1781, trans. N. Kemp Smith. New York: Macmillan Press.

Kim, Jaegwon, 1973, "Causation, Nomic Subsumption, and the Concept of Event," *Journal of Philosophy*, 70: 217–36.

——, 1976, "Events as Property Exemplifications," in M. Brand and D. Walton (eds.), *Action Theory*, Dordrecht: D. Reidel Publishing, pp. 159–77.

Kistler, Max, 1998, "Reducing Causality to Transmission," *Erkenntnis*, 48: 1–24.

Kutach, Douglas, 2007, "The Physical Foundations of Causation," in H. Price and R. Corry (eds.), *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, Oxford: Oxford University Press, pp. 327–50.

Kvart, Igal, 1986, *A Theory of Counterfactuals*, Indianapolis: Hackett Publishing.

——, 1997, "Cause and Some Positive Causal Impact," in J. Tomberlin (ed.), *Philosophical Perspectives 11: Mind, Causation, and World*, Oxford: Basil Blackwell, pp. 401–32.

——, 2004, "Probabilistic Cause, Edge Conditions, Late Preemption and Discrete Cases," in P. Dowe and P. Noordhof (eds.), *Cause and Chance: Causation in an Indeterministic World*, London: Routledge.

Lewis, David, 1979, "Counterfactual Dependence and Time's Arrow," *Noûs*, 13: 455–76.

——, 1986a, "Causation," in *Philosophical Papers 2*, Oxford: Oxford University Press, pp. 159–213.

——, 1986b, "Events," in *Philosophical Papers 2*, Oxford: Oxford University Press, pp. 241–69.

——, 2000, "Causation as Influence," *Journal of Philosophy*, 97: 182–97.

——, 2004, "Void and Object," in J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*, Cambridge, MA: The MIT Press, pp. 277–90.

Loewer, Barry, 2007, "Counterfactuals and the Second Law," in H. Price and R. Corry (eds.), *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, Oxford: Oxford University Press, pp. 293–326.

Mackie, J. L., 1965, "Causes and Conditions," *American Philosophical Quarterly*, 2: 245–64.

——, 1974, *The Cement of the Universe*, Oxford: Oxford University Press.

Maslen, Cei, 2004, "Causes, Contrasts, and the Nontransitivity of Causation," in J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*, Cambridge, MA: The MIT Press, pp. 341–58.

Maudlin, Tim, 2007, "On the Passing of Time," in *The Metaphysics within Physics*, Oxford: Oxford University Press, pp. 104–42.

McDermott, Michael, 1995, "Redundant Causation," *British Journal for the Philosophy of Science*, 46: 423–44.

——, 2002, "Causation, Influence, and Sufficiency," *Journal of Philosophy*, 99: 84–101.

McGrath, Sarah, 2005, "Causation by Omission: A Dilemma," *Philosophical Studies*, 123: 125–48.

Mellor, D. H., 1981, *Real Time*, Cambridge: Cambridge University Press.

——, 1988, "On Raising the Chances of Effects," in J. Fetzer (ed.), *Probability and Causality*, Dordrecht: D. Reidel Publishing, pp. 229–40.

——, 1995, *The Facts of Causation*, London: Routledge Press.

Menzies, Peter, 1989a, "A Unified Account of Causal Relata,"

Australasian Journal of Philosophy, 67: 59–83.

—, 1989b, "Probabilistic Causation and Causal Processes: A Critique of Lewis," *Philosophy of Science*, 56: 642–63.

—, 1996, "Probabilistic Causation and the Pre-emption Problem," *Mind*, 105: 85-117.

—, 2004, "Difference-Making in Context," in J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*, Cambridge, MA: The MIT Press, pp. 139–80.

—, 2007, "Causation in Context," in H. Price and R. Corry (eds.), *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, Oxford: Oxford University Press, pp. 191–223.

Menzies, Peter and Huw Price, 1993, "Causation as a Secondary Quality," *British Journal for the Philosophy of Science*, 44: 187–203.

Mill, J. S., 1846, *A System of Logic*, New York: Harper & Brothers.

Moore, Michael, 2009, *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics*, Oxford: Oxford University Press.

Noordhof, Paul, 1999, "Probabilistic Causation, Preemption and Counterfactuals," *Mind*, 108: 95–125.

Northcott, Robert, 2008, "Causation and Contrast Classes," *Philosophical Studies*, 139: 111–23.

Papineau, David, 1993, "Can We Reduce Causal Direction to Probabilities?," in Hull, D. M. Forbes, and K. Okruhlik (eds.), *PSA 1992* (Volume 2), East Lansing: Philosophy of Science Association, pp. 238–52.

Paul, L. A., 1999, "Keeping Track of the Time: Emending the Counterfactual Analysis of Causation," *Analysis*, 58: 191–8.

—, 2000, "Aspect Causation," *Journal of Philosophy*, 97: 223–34.

Paul, L. A. and Ned Hall, 2013, *Causation: A User's Guide*, Oxford: Oxford University Press.

Pearl, Judea, 2000, *Causality*, Cambridge: Cambridge University Press.

Price, Huw, 1991, "Agency and Probabilistic Causality," *British Journal for the Philosophy of Science*, 42: 157–76.

—, 1996, *Time's Arrow and Archimedes' Point*, Oxford: Oxford University Press.

—, 2007, "Causal Perspectivalism," in H. Price and R. Corry (eds.), *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, Oxford: Oxford University Press, pp. 250–92.

Quine, W. V. O., 1966, *The Ways of Paradox*, New York: Random House.

—, 1985, "Events and Reification," in E. LePore and B. McLaughlin (eds.), *Actions and Events: Essays on the Philosophy of Donald Davidson*, Oxford: Basil Blackwell, pp. 162–71.

Ramachandran, Murali, 1997, "A Counterfactual Analysis of Causation," *Mind*, 106: 604–13.

Russell, Bertrand, 1992, "On the Notion of Cause," orig. 1912, in J. Slater (ed.), *The Collected Papers of Bertrand Russell v6: Logical and Philosophical Papers 1909–1913*, London: Routledge Press, pp. 193–210.

—, 1948, *Human Knowledge: Its Scope and Limits*, New York: Simon and Schuster.

Salmon, Wesley, 1984, *Scientific Explanation and the Causal Structure of the World*, Princeton: Princeton University Press.

—, 1997, "Causality and Explanation: A Reply to Two Critiques," *Philosophy of Science*, 64: 461–77.

—, 1998, *Causality and Explanation*, Oxford: Oxford University Press.

Sartorio, Carolina, 2005, "Causes as Difference-Makers," *Philosophical Studies*, 123: 71-96.

Sayre, Kenneth, 1977, "Statistical Models of Causal Relations," *Philosophy of Science*, 44: 203–14.

Schaffer, Jonathan, 2000a, "Trumping Preemption," *Journal of Philosophy*, 97: 165–81.

—, 2000b, "Overlappings: Probability-Raising without Causation," *Australasian Journal of Philosophy*, 78: 40–6.

——, 2000c, "Causation by Disconnection," *Philosophy of Science*, 67: 285–300.

——, 2001, "Causes as Probability-Raisers of Processes," *Journal of Philosophy*, 98: 75–92.

——, 2005, "Contrastive Causation," *Philosophical Review*, 114: 327–58.

——, 2007, "Review: Dowe and Noordhof's *Cause and Chance*," *The British Journal for the Philosophy of Science*, 58: 869–74.

——, 2012, "Causal Contextualisms," in M. Blaauw (ed.), *Contrastivism in Philosophy*, London: Routledge, pp. 35–63.

Shoemaker, Sydney, 1980, "Causality and Properties," in P. van Inwagen (ed.), *Time and Cause: Essays Presented to Richard Taylor*, Dordrecht, D. Reidel Publishing, pp. 109–35.

——, 1998, "Causal and Metaphysical Necessity," *Pacific Philosophical Quarterly*, 79: 59–77.

Skyrms, Brian, 1984, "EPR: Lessons for Metaphysics," in P. French, T. Uehling, Jr., and H. Wettstein (eds.), *Midwest Studies in Philosophy IX*, Minneapolis: University of Minnesota Press, pp. 245–55.

Sober, Elliott, 1985, "Two Concepts of Cause," in P. Asquith and P. Kitcher (eds.), *PSA 1984* (Volume 2), East Lansing: Philosophy of Science Association, pp. 405–24.

Spirtes, Peter, Clark Glymour, and Richard Scheines, 1993, *Causation, Prediction, and Search*, New York: Springer-Verlaag.

Strawson, P. F., 1985, "Causality and Explanation," in B. Vermazen and M. Hintikka (eds.), *Essays on Davidson: Actions and Events*, Oxford: Clarendon Press, pp. 115–36.

Steward, Helen, 1997, *The Ontology of Mind*, Oxford: Clarendon Press.

Suppes, Patrick, 1970, *A Probabilistic Theory of Causality*, Amsterdam: North Holland Publishing.

Swain, Marshall, 1978, "A Counterfactual Analysis of Event Causation," *Philosophical Studies*, 34: 1–19.

Taylor, Richard, 1966, *Action and Purpose*, Upper Saddle River: Prentice Hall.

Tooley, Michael, 1987, *Causation: A Realist Approach*, Oxford: Clarendon Press.

——, 2004, "Probability and causation," in P. Dowe and P. Noordhof (eds.), *Cause and Chance: Causation in an Indeterministic World*, London: Routledge, pp. 77–119.

Van Fraassen, Bas, 1980, *The Scientific Image*, Oxford: Oxford University Press.

Vendler, Zeno, 1984, "Agency and Causation," in *Midwest Studies in Philosophy*, 9: 371–84.

Von Wright, G. H., 1975, "On the Logic and Epistemology of the Causal Relation," in *Causation and Conditionals*, Oxford: Oxford University Press, pp. 95–113.

Weslake, Brad, forthcoming, "A Partial Theory of Actual Causation," *The British Journal for the Philosophy of Science*,

Woodward, James, 1984, "A Theory of Singular Causal Explanation," *Erkenntnis*, 21: 231–62.

A Theory of Explanation, 2003, *A Theory of Explanation*, Oxford: Oxford University Press.

Yablo, Stephen, 2002, "De Facto Dependence," *Journal of Philosophy*, 99: 130–48.

Academic Tools

⚘ How to cite this entry.

⚘ Preview the PDF version of this entry at the Friends of the SEP Society.

⚘ Look up this entry topic at the Indiana Philosophy Ontology Project (InPhO).

PP Enhanced bibliography for this entry at PhilPapers, with links to its database.

Other Internet Resources

- Causality, entry in Wikipedia.
- Causation, Laws, etc., category at PhilPapers.

Related Entries

causation: and manipulability | causation: backward | causation: causal processes | causation: counterfactual theories of | causation: in the law | causation: probabilistic | conditionals: counterfactual | decision theory: causal | determinism: causal | events | Hume, David | laws of nature | mental causation | metaphysics | physics: Reichenbach's common cause principle | probability, interpretations of | properties | time | time travel | time travel: and modern physics | tropes

Acknowledgments

Thanks to Phillip Bricker, Christopher Read Hitchcock, Doug Kutach, Laurie Paul, and Ignacio Silva.