

SPECIFICITY AND REDUNDANT CAUSATION

HENNING STRANDIN

ABSTRACT. In this paper I present a metaphysically minimalist but theoretically strong version of fact causation, in which the causal relata constitute a full Boolean algebra, mirroring the entailment relation of the sentences that express them. I suggest a generalization of the notion of multiple realizability of causes in terms of specificity of facts, and employ this in an interpretation of what goes on in cases of apparently redundant causation.

1. CAUSES, PARTICULARS, AND FACTS

I think that what exactly the causal relation relates is a very hard question. A common intuition is that, since causation seems to happen in the world, the causal relata must exist in the world. Events, objects, and instances of properties exist in the world, while propositions and facts, as ordinarily conceived, do not, but are about the world. However, in *most* of our actual causal talk we don't explicitly designate events or objects or property instances as causes. We say things like

A: When the electromagnet was turned on, it caused the electron to emit a photon.

B: The train was late because there was something wrong with one of the cars.

C: That I dislike *Julius Caesar* is a consequence of my having to study it in high school.

C is from Anthony Bennett's (Bennett 1988, ch. II). In that chapter, he also explains why **A**, which is in what he calls noun-infinitive form, must be relating facts. To

translate his reasoning to our example, while turning on the electromagnet caused the electron to emit a photon, it is not true that

A': When the electromagnet was turned on, it caused the electron to emit a photon at t .

Turning on the magnet did not cause the electron to emit a photon at the time t , when it was emitted, since the time of the emission was not determined by law but a stochastic aspect of the event. Turning on the magnetic field was not sufficient for the emission to happen at t —nothing was, as far as we know. **A** and **A'** “name different facts, one richer than the other, so they could differ in truth-values” (Bennett 1988, p. 25). However, the event of the photon emission was an event that took place at t , so if making a true causal claim is a matter of referring to the right event, then “the photon emission” and “the photon emission that occurred at t ” must refer to the same event and an event-relating causal claim using the first term cannot differ in truth value from one using the second. This observation foreshadows the main theme in this paper. It is also reminiscent of Laurie Paul’s criticism of the event account of causation, and her suggestion that it is event *aspects* that are causally related (Paul 2004).

All of **A**, **B**, and **C** relate facts as causes, then. We do of course also sometimes explicitly relate things in the world, causally. As in “the bump in the road was the cause of the damage to the car.” This designates the bump in the road as the cause. But we know that it is very unlikely that physical objects can be the causes in any analysis of causation. So even if that formulation is a perfectly acceptable and intelligible causal claim, what we *really* mean must be that the bump’s presence in the road, or the fact that the bump was in the road, was the cause, rather than just the bump. Both “the presence of the bump caused the damage to the car” and “the car was damaged because there was a bump in the road” do a better job of articulating the cause. I take it that the former refers to an event as the cause, and the latter relates facts. They don’t say exactly the same thing, of course. The former

has a singular term referring to a particular event (the presence of the bump), while the latter claim could better be understood as having an existential generalization, true if at least one bump was present in the road. And there is no *competition* here, between referring by a singular term and expressing a fact. I take it that just as we can refer to Mozart's death, we can refer to the fact that Mozart's death occurred in 1791 (like that). When we say "Mozart died in 1791" we are strictly speaking expressing a different fact—the fact that Mozart died at least once that year. It's something less specific, but it works well enough as cause or effect in many causal claims. "Because Mozart died in 1791, he did not finished the Requiem."

Fact causes, then, can involve particulars of every kind, for example when expressed by sentences containing a reference or an existential generalization, and this seems independent of the surface form of a causal claim. The converse—that true causal claims always relate particulars—is more complicated. True causal claims can designate absences and perhaps (as I will argue) disjunctive things as causes, and it is hard to see how those always correspond to some actual particular such that it could be causally related. I think (Schaffer 2004) ought to convince us that absence causes are normal causes. In that article Schaffer stacks examples of paradigmatic cases of causation and uncontroversial causal claims from science and law, that all involve absence causation (usually in the form of double prevention), and he also discusses adequacy criteria on theories of causation in light of theses things. For now I ask you to assume—for the sake of argument—that when particulars are involved in singular causation they are so by being involved in the relevant facts. I will talk of these facts as causally related, even if that relation, which relates facts that may involve objects as well as times and places, turns out to be different in kind from one relating occurrent objects that exist in the world *at* times and places.

The fact that is expressed by a true contingent sentence I take to be the truth-maker of that sentence—and the state of affairs expressed by a possibly true contingent sentence the possibility-maker of that sentence. I will not claim that there is such a thing making every true sentence true; I'm interested only in things that plausibly can occur in possibly true causal claims, and that's contingent things. Mere possibilities are different from actualities, but even a mere possibility can make a sentence merely possibly true. If there are possibilities, that is—and I think there are. A deterministic physical law that allows for different initial conditions of a system allows for different time evolutions of the system, and the set of such time evolutions is the space of possibilities that is implied by that law. (An indeterministic law need not even allow different initial conditions to imply an interesting—albeit not a very causally interesting—possibility space.) I think the world is governed by some law, that is necessarily underdetermined by the actual history of the world, and thus that the world is uniquely determined only by its actual history together with the law that governs it, and that this law can be understood as a space of possibilities. In short, the world is all the actualities and the possibilities. But not every possibility is in accordance with the actual laws of nature. Possibilities seem to form a hierarchy. That I am a married bachelor is a mere logical possibility, and that I fly to work in the morning by flapping my arms is a mere conceptual possibility. Neither is, I take it, a nomic possibility, and these are the most interesting ones when considering causes (even if we also importantly indulge in counter-legal causal reasoning at times), since causations are usually believed to be law-governed.

I'll now suggest a metaphysically “thin” theory for states of affairs and facts, in which they are modeled as sets of possible worlds. Don't think of these sets of possible worlds as course-grained propositions. Think of them rather as *world properties*. Something like this: take “The sentence ϕ is true in $_$ ” to express a world property, the property that makes ϕ true, and that is interpreted as a set

of possible worlds much like a predicate is interpreted as the set of objects that have the property expressed by the predicate. A discussion about whether things expressed must be things we can understand, whether truthmakers can be things imbued with meaning, or, if not, where exactly understanding ends and truthmaking begins, may be important but is way beyond the scope of this text. Likewise, there may be sentences with different meanings that are true in exactly the same worlds (different tautological sentences, for example), but if so, then that difference does anyway not correspond to a difference between two worlds in our possibility space, and we are interested mainly in such differences when investigating possible causes and effects. There, then, is where we should begin.

2. DEFINING STATES OF AFFAIRS

If a sentence differs in truth value between two possible worlds, then these two worlds are in some respect different. The power set $P(W)$ of a set W of possible worlds contains all possible ways of sorting the worlds based on some difference between them. That is, for any meaningful contingent sentences “ ϕ ” and “ $\neg\phi$ ” there are non-empty sets s and $W \setminus s$ (the complement set of s in W) such that the worlds in s differ from the worlds in $W \setminus s$ exactly in the way that makes “ ϕ ” true rather than false. More precisely, let \mathcal{L} be some alethic language for which we already have a possible-world semantic theory, so that we have for \mathcal{L} the relation of semantic truth of a sentence in a world w , $\models_w \phi$, and of semantic entailment, $\Sigma \models \phi$, where Σ is a set of sentences. (The relation is of course also relative to a model and an assignment of values to any variables, but this will remain implicit in our symbolic expressions.) Let W be the set of possible worlds in that theory, and C the set of contingent sentences in \mathcal{L} . (We require that a sentence expressing a cause or an effect is contingent, so won’t assign states of affairs to any other.) We can then define a function $S : C \rightarrow P(W)$ that assigns states of affairs to contingent sentences:

Definition 1. $S(\phi) = \{w \in W \mid \models_w \phi\}$.

That is, the state of affairs expressed by ϕ is interpreted in the theory as the set of worlds in which ϕ is true. The set of states of affairs constitute, as expected, a Boolean algebra:

- (1) $S(\neg\phi) = W \setminus S(\phi)$,
- (2) $S(\phi \vee \chi) = S(\phi) \cup S(\chi)$, and
- (3) $\Sigma \models \phi$ iff $\bigcap_{\chi \in \Sigma} S(\chi) \subseteq S(\phi)$.

One thing worth mentioning is that this theory is immune to the slingshot argument against fact causation and any interesting correspondence between sentences and facts in general. As employed by Donald Davidson in (Davidson 1967b), the slingshot proves—given certain seemingly plausible semantic principles—that if any fact is a cause of something, then every fact is a cause of that thing. The principles on which the proof relies are substitution *salva veritate* of logically equivalent alternatives to ϕ and χ in a causal claim “ χ because ϕ ,” and also of co-extensional singular terms inside ϕ and χ . We will not dispute the first substitution principle. Let the claim be “the pig was roasted because there was a fire in Jones’s house.” Now Davidson says that,

if the fact that there was a fire in Jones’s house caused it to be the case that the pig was roasted, and Jones’s house is the oldest building on Elm street, then the fact that there was a fire in the oldest building on Elm street caused it to be the case that the pig was roasted. We must accept the principle of extensional substitution, then. (Davidson 1967b, p. 694.)

This substitution relies on that “Jones’s house” and “the oldest building on Elm street” refer to the same object. That may be the case, and still the fact that there was a fire in Jones’s house should not be identical to the fact that there was a fire in the oldest building on Elm street. And for us, it clearly isn’t. There are

possible worlds in which Jones's house is not the oldest building on Elm street, so the two sentences are not true in exactly the same worlds, hence they do not express the same fact. That is, substitution of co-extensional singular terms does not go through in general. More straightforwardly, since our theory already has an identity criterion for facts (the axiom of extensionality, given that they are sets), no language for which our theory is an interpretation, and that has singular terms denoting facts, can prove the slingshot and be sound. (See Restall 2004 for a nice paper on this, with proof.) With the basics in place, we can look at two issues that are related: the multiple realizability of causes; and that of absence causation already alluded to.

3. REALIZERS AND ABSENCES

Imagine a pigeon, Sophie, conditioned to peck at red to the exclusion of other colors; a red triangle is presented, and Sophie pecks. Most people would say that the redness was causally relevant to her pecking, even that this was a paradigm case of causal relevance. But wait! I forgot to mention that the triangle in question was a specific shade of red: scarlet. Assuming that the scarlet was causally sufficient for the pecking, we can conclude by the exclusion principle that every other property was irrelevant. Apparently, then, the redness, although it looked to be precisely what Sophie was responding to, makes in reality no causal contribution whatever. (Yablo 1992, p. 257.)

The roles in causation of multiply realizable causes and their realizers have been discussed in particular in the literature about mental causation. Broadly speaking, the common belief that mental properties, states, or events supervene on physical ones has seemed to introduce a dilemma. We often take a mental event, such as my feeling peckish, to be the cause of some action, such as my grabbing a snack.

But if some physical event is what realizes my peckishness, and it moreover is causally sufficient under the conditions on the occasion for the physical event that realizes my snack-grabbing, then either there is no causing left for the mental event to do (by the exclusion principle, mentioned by Yablo above), or we are faced with a case of causal redundancy. It has frequently been recognized that the problem generalizes to cases outside mental causation, as Yablo's example is meant to illustrate. Jaegwon Kim compares the mental/physical distinction to that between macrophysical and microphysical. His account is reductionist, but he does not think that reducibility warrants elimination, but rather that it validates that which has been shown reducible, at least in the case of the mental. Thus, he thinks that mental causation "does take place; it is only that it is epiphenomenal causation, that is, a causal relation that is reducible to, or explainable by, the causal process taking place at a more basic physical level" (Kim 1993, p. 107). He moreover believes that "all causal relations involving observable phenomena—all causal relations familiar from daily experience—are cases of epiphenomenal causation" (Kim 1993, p. 95). Yablo frames his account in terms of determinables and their determinates. He, too, wants to explain how mental events can be determined by physical events but how they nevertheless are neither causally inert nor produce systematic causal overdetermination. He appeals to a notion of causes being "commensurate" to their effects: "roughly, they should incorporate a good deal of causally important material but not too much that is causally unimportant" (Yablo 1992, p. 274), and he also calls this "proportionality." He observes that sometimes a suggested cause is "overly determined," that is to say, it is too specific and therefore not necessary for the effect in question—and sometimes it not determined enough, too general, and therefore fails to be sufficient for the effect. The scarlet color of the triangle is too specific to be the cause of Sophie's pecking, given that she was trained to pick at red things. The pecking would have occurred even if the triangle had been some different shade of red. Claiming, on the other hand, that it was that the triangle

was colored that was the cause is too general, the triangle could have been colored without it causing Sophie to peck at it. Yablo recognizes that a cause that is too determined or not sufficiently determined for the effect at hand can have a *different* effect, to which it is more proportional.

Then when do we attribute effects to mental causes? Only when we believe, I can only suppose rightly, that the effect is relatively insensitive to the finer details of [the mental event] *m*'s physical implementation. Having decided to push the button, I do so, and the doorbell rings. Most people would say, and I agree, that my decision had the ringing as one of its effects. Of course, the decision had a physical determination *p*; but, most people would also say, and I agree again, that it would still have been succeeded by the ringing, if it had occurred in a different physical way, that is, if its physical determination had been not *p* but some other physical event. And this is just to say that *p* was not required for the effect. (Yablo 1992, p. 278.)

All I want to add to this picture is the idea that the physical realizer *p* likely had an effect, namely a realizer of the bell ringing, such that *p* was required for *that* specific effect. That is, the fact of the occurrence of the mental event, being such that it could be realized by one of several facts about physical events, is, as I will call it, less specific than its physical realizer, and therefore it can at least sometimes be expected to have an effect that is distinct from that of its realizer, in that the effect of the mental event is proportionally general, in some loose and intuitive sense. (This language has precedence in (Gibbons 2006, p. 94). In regards to Sophie's pecking, '[r]ed wins [over scarlet] without relying on magic because it is at the right level of generality. It's not too specific; it's not too general; it is, as Stephen Yablo says, "proportional" to the effect.'))

I want to suggest that these impressions of causes being proportional to their effects, and of—on one and the same occasion—a more general cause having a more general effect, is directly related to the observation by Bennett quoted above, regarding the difference in “richness” (I will call it specificity) between different facts obtaining on the occasion of one and the same causation. I also agree that on pretty much any occasion for which some true causal claim is made, a more specific cause can be denoted or expressed, that will have a more specific effect. That is to say, that cause is more specific than the one first claimed, and the effect of *that* cause will have the same relation to the originally stated effect. My feeling peckish on some occasion is realized by some particular physical state in my nervous system. It, along with other necessary conditions, relegated to the background for pragmatic reasons, caused a particular movement of my body that qualifies as a snack-grabbing. If my peckishness really caused my snack-grabbing under those conditions, then if it had been realized by a different physical state, then that state might, probably, have caused a different movement of my body, but that movement would also have qualified as a snack-grabbing. I take this to be what it means to say that it *was* my peckishness that caused my snack-grabbing. If some other realizer of that mental state, possible under the relevant conditions, had not caused me to grab a snack, then the mental state was not sufficient for the purported effect. (This ignores the possibility of causes that merely raise the probability of the effect. Throughout this paper I will assume that causes are both necessary and sufficient for their effects under those relevant conditions that obtained on the occasion of the causation. It is easier to maintain this condition on causes if we can reinterpret cases of apparently redundant causation, which is just what I try to do in this paper.)

This view is clearly related to the discussion about fragility and robustness of events in event-relation theories of causation. (See, for example, Lewis 1987, Appendix E.) I will return to the classical examples of causal redundancy shortly, let

me just point out now that the problems in the event theories of causation, of making events too fragile or too robust, are not mirrored in the fact account. We never need to choose a level of specificity of facts—all the facts are there at all times—and they are not indiscriminately more or less fragile, but just in the way expressed by a corresponding sentence.

The facts of causation, it is then suggested, are usually multiply realizable. And this is of course why they are modeled as sets of possible worlds. Each element in that set is a possible realizer for that fact. Not only that, each subset of that set is a fact, and a realizer of any fact that is a superset of it. There are many more facts than there are sentences, but we knew that already.

Moreover, if we accept that facts, understood as properties of worlds and modeled as sets of possible worlds, can be causes, and that these are moreover multiply realizable when they are sets of more than one possible world, then absence causes need not appear very mysterious. I'm wearing green shoes. This fact is multiply realizable. My shoes have a certain shade of green, but the claim would still have been true, and would still have expressed the same fact, if they had been a different shade. This is also true of the possible state of affairs of my not wearing green shoes. There are all kinds of ways in which I can wear green shoes, and in which I can fail to do so. The fact of my not wearing green shoes is perhaps more "disparate," in some sense. I guess not-green things is not a natural kind. But whether they are of the *appropriate* kind for causal claims ought reasonably be determined by the possible truth of causal claims in which they are involved. If we believe that absences can be causes, then it is at least somewhat comforting to note that they are not essentially different kinds of things from presences, in that they are usually multiply realizable states of affairs, expressible by sentences.

If we have some theory of specificity of sentences, we also have a theory of specificity of facts, and to this we now turn.

4. RELATIVE SPECIFICITY

There are at least a couple of ways to increase or decrease the specificity of a sentence. We can, for example, add and remove predicates, adverbs and adverbial phrases in a sentence such as

ϕ : Doris sneezed suddenly in the cold bedroom at [8 : 16 – 8 : 17].

We take the time to be expressed by a minute interval, for reasons that become clear below. Davidson has shown us one logical form we can give to such a sentence if we want to preserve the intuitive entailment from ϕ to

χ : Doris sneezed in the bedroom.

(Davidson 1967a, and see Ludwig 2010 for later developments and refinements.) I will take it as a given that ϕ is a more specific statement, in the sense of providing more information about what is going on, than χ is. The sentence reasonably states that the bedroom is cold when Doris sneezes, and not when the sentence is evaluated, so, to get access to the structure of the predicate “cold” as well as of the adverb “suddenly” we analyze “cold” as an event, the bedroom *being* cold, occurring at the same time as the sneezing. Taking e, e' to be variables ranging over events, and t to range over minute intervals, we can then analyze ϕ as

$$\phi\text{-f: } \exists e \exists e' \exists t (sneezing(e) \wedge by(Doris, e) \wedge sudden(e) \wedge in(the\ bedroom, e) \wedge at(t, e) \wedge \\ being\ cold(the\ bedroom, e') \wedge at(t, e') \wedge t = [8 : 16 - 8 : 17])$$

$\phi\text{-f}$ entails $\chi\text{-f}$ by \wedge -elimination.

$$\chi\text{-f: } \exists e (sneezing(e) \wedge by(Doris, e) \wedge in(the\ bedroom, e))$$

Unsurprisingly then, a more specific account of some goings on entails a less specific account of the same goings on. Our notion of specificity is very weak. There is of course no *measure* of specificity (although, had we had reason to think the possibility space finite, cardinality might have been a candidate), and we can't even compare the relative specificity between arbitrary sentences. Specificity follows logical strength. At least for contingent sentences. Depending on what we take

the *meaning* of a tautology such as “it rains or it doesn’t” to be, we may or may not be willing to say that it is less specific than “Doris sneezed in the bedroom,” even though the latter entails the former. But, surely, a contradiction is not the most specific sentence relative to anything? Regardless, we are only concerned with contingent sentences.

Adding and removing predicates, adverbs, and adverbials is a way of changing the specificity of a sentence that is close to how common language works. There is another way of treating exactly the same situation, that is farther from common language but in a way more easy to think about. (Your mileage may vary.) We can just put the different contingent possibilities that are compatible with our claim in disjunction. Say that we don’t know when Doris sneezed. “Doris sneezed in the bedroom” is equivalent to (again using minute intervals)

ψ : Doris sneezed in the bedroom at [00 : 00 – 00 : 01) or Doris sneezed in the bedroom at [00 : 01 – 00 : 02) or ... or Doris sneezed in the bedroom at [23 : 59 – 00 : 00).

Clearly, ϕ entails ψ , since it entails χ . We can then increase specificity by deleting disjuncts. Adding the adverbial “at [8 : 16 – 8 : 17)” to χ is equivalent to deleting all disjuncts except the one stating that the sneezing happened at [8 : 16 – 8 : 17) from ψ —and results in the same sentence. Of course, this doesn’t tell us when in the interval [8 : 16 – 8 : 17) the sneezing occurred, so the resulting sentence is still equivalent to a disjunction of sentences giving the time with, for example, second precision. Rinse and repeat. From the point of view of the facts, the one expressed by “Doris sneezed in the bedroom” can be partitioned in many different ways, and some of those ways will correspond to some particular aspect of how that fact could be realized, such as its time. In many such cases, each part is expressible by a more specific sentence, such as one of the disjuncts in ψ . What we are after here is mainly the idea that, in all likelihood, nearly *any* sentence plausibly expressing a causal fact is equivalent to a disjunction of sentences expressing more specific

possible ways for that fact to be realized. And, since specificity follows entailment, *any* contingent disjunction such that it entails none of its disjuncts expresses a less specific fact than either of them. This is a broad generalization of the notion of multiple realizability of causes described in the previous section.

5. REDUNDANT CAUSATION

5.1. **Causes are necessary on the occasion.** It is a presumption in talk about redundant causation that the cause of some effect on a particular occasion (a “singular cause”) is necessary for the effect under the circumstances that obtained on that occasion. That is, given the way things actually were, had the cause not happened neither would the effect. If this expectation is relaxed, then the problem of redundancy is diminished to the same degree. I will assume that causes are, without exception, necessary on the occasion, and I take that to be the strongest intuition we have about causes. This does not mean that I don’t believe in general causes that increase the probability of a general effect, but I resist analyses of causation that take probabilities as fundamental. Partly this is because it is difficult enough to understand what the *probabilities* are, so that they don’t necessarily contribute to our understanding of causes. I will assume that the world is near-enough deterministic on the macroscopic scale, and that probabilities belong to general causation and to statistical methods for finding causes. If smoking causes cancer, then on some particular occasion it either was a necessary condition for that cancer, or it wasn’t a cause of it. The probabilities enter because we don’t know what other conditions are necessary in general together with smoking for cancer to occur, or exactly when these conditions obtain. There are likely also truly stochastic facts, such as the precise time of some quantum event. This kind of indeterminism can certainly propagate to the macroscopic world, for example if we were to use a true random number generator to control some process. However, since we can distinguish facts finely—as in making a distinction between the electron emitting a photon and

it emitting a photon at t —we can allow for that these facts are not caused at all (since they are not even determined by natural law). Thus, causes are generally necessary for their effects on the occasion. What seems to have most effectively convinced people otherwise is examples of apparently redundant causation, and we turn to those now.

5.2. Symmetric overdetermination and event fragility. For some perfectly possible causal scenarios, it seems to many that there is nothing that both is of a kind such that it could be a cause and is necessary for the effect under the circumstances. Rather, several cause-like factors are each sufficient for the effect on that occasion. There are several distinct kinds of causal redundancy in the literature, mainly the cases of preemption and of symmetrical overdetermination. My treatment is similar for both. I will start with Lewis’s characterization of what has come to be called symmetrical overdetermination. (I think it is Lewis’s term originally, although he calls it “redundancy” in this passage.)

Suppose we have two events c_1 and c_2 , and another event e distinct from both of them; and in actuality all three occur; and if either one of c_1 and c_2 had occurred without the other, then also e would have occurred; but if neither c_1 nor c_2 had occurred, then e would not have occurred. Then I shall say that c_1 and c_2 are redundant causes of e . (Lewis 1987, p. 31.)

In this passage, Lewis describes what could be understood as a *disjunctive cause* of e . It is the disjunction of the events c_1 and c_2 . Such disjunctions have appeared before, and after, in the philosophical literature on causation. In (Mackie 1965, p. 26), J. L. Mackie says about a possible case of the same general form, “In such cases, the requirements of my analysis [...] are not met: each proposed cause *is* redundant and not even necessary *post factum*, though the disjunction of them is necessary *post factum* and non-redundant.” Both Mackie and Lewis take the causal

relata to be events, and they both reject the disjunction as being the cause in these examples. And both appeal to similar explanations for why these cases fail in their theories without making the theories inadequate: on such occasions when both c_1 and c_2 actually happened and each was sufficient for e , our everyday judgment as to what was the cause would be unclear (Mackie 1965, p. 26, and Lewis 1987, p. 31). The theory is adequate, on this particular point, in virtue of following everyday intuitions.

Lewis furthermore observes that whether there is a case of redundancy in any realistic example of this kind depends on how finely we individuate the events in question, in the possibility space. That is, how different can the *same* event be as it occurs in a different possible world? Let's take a killing as an example. (It is a tradition.) We want to say that A 's shooting B was the cause of B 's death. Lewis remarks that just because B would have died at some point anyway, so A 's shooting B was not necessary for B 's *eventual* death, that doesn't make this a case of redundant causation. It's only a case of redundancy if some other event that actually occurred was sufficient for bringing about the *same* death as the one that happened. Similarly, if A and C both shot B simultaneously, it is still only a case of redundant causation if the death that occurs when both A and C fire is the same as the one that occurs when only A fires, and also the same as the one that occurs when only B fires. In the first example, when only A fires, it is clearly unreasonable to take B 's eventual death, that would undoubtedly have happened if A had not shot B , but possible many years later, as the same death as the one caused by A 's shooting. But in the last example it may seem unreasonable to say that the deaths occurring when both A and B fire, when only A fires, and when only B fires, are distinct events. How fragile or robust we take the events involved in the causation to be makes all the difference.

Lewis elaborates further on the problem of taking events to be very fragile:

[S]uppose there was a gentle soldier on the firing squad, and he did not shoot. If the minute difference made by eight bullets instead of seven is enough to make a different event, then so is the minute difference made by eight instead of nine. So if the victim's death is so very fragile that it would not have occurred without your act, equally it is so fragile that it would not have occurred without the gentle soldier's omission. If by reason of fragility the death depends causally on your act, then equally it depends causally on the omission. So the gentle soldier caused the death by not shooting, quite as much as you caused it by shooting! (Lewis 1987, p. 35.)

The general idea is that if we make the event in a case of apparent redundant causation so fragile that we can distinguish a uniquely sufficient cause, then it also becomes dependent on things we intuitively would never call causes of that event. But, to complicate things further (this complication is mine and not Lewis's), we can supposedly refer to the same death by the term "the eight-bullet death" (since it was in fact an eight-bullet death). If the rules for executions state that every execution should be an at-least nine-bullet death, then our gentle soldier might be in some trouble by having caused the rules to be violated. It's the same event, with supposedly the same causes, so why is this causal claim suddenly more reasonable? Perhaps we *adjust* the modal fragility of events from one claim to another, based on context and pragmatic concerns? Lewis suggests that something like a "double standard" is involved in such cases as when we on the one hand want to reject the gentle soldier's not shooting as a cause of the death of the prisoner, and on the other hand accept his not shooting as a cause of the eight-bullet death of the prisoner. If we think that an event cannot both have an event as its cause and not have that event as its cause, then "the death" and "the eight-bullet death" must refer to different events. But they are certainly related, because it was that death

that was an eight-bullet death. Lewis is not suggesting that he has a solution to this problem.

I have emphasized two features in this discussion about symmetrically overdetermined causation. One is the observation, made early in the literature, that in these cases a disjunction does satisfy the necessity condition on causes—the other is that how modally fragile we take events to be can have dramatic effects on the intuitive truth values of causal claims involving those effects. Before I go on to how these things connect to our understanding of causes as facts and our interpretations of disjunctive sentences and degrees of specificity, I should say something about the other major kind of causal redundancy: preemption. But even before that I should mention that Hugh Mellor has also made the observation that disjunctions qualify as causes in cases of redundancy, and that he, advocating facts as causes but facts substantially different from ours, embraces them.

[I]n the circumstances S , if Don had neither fallen [C] nor been shot [C'], his chance of dying [E] would have been zero. So between them C and C' do provide one deterministic cause of E , namely $C \vee C'$: Don dies because he falls *or* is shot. (Mellor 1995, p. 102.)

Mellor defends this claim but does not provide an intuitive understanding of the cause in question. I don't know how sympathetic he would be to the proposal I make in this text.

5.3. Preemption, early and late. Imagine a case such as the one outlined by Lewis in the first quote above, but in which c_2 does not in fact happen, and where c_2 *would* have happened if c_1 had not. That is, c_1 is a *preventer* of c_2 , and both c_1 and c_2 are individually sufficient for the effect e on the occasion. An old example case is from Michael Scriven, retold by Mackie (Mackie 1965, p. 25): a man suffers a heart attack at t_1 , which causes him to die at t_3 , although, had he

not had the heart attack, he would have had a stroke at t_2 , again killing him at t_3 . There is broad agreement that in a case such as this, common intuitions strongly suggest that it is c_1 that is the cause. After all, c_2 didn't even happen. "It is clear what answer we want—the preempting cause is a cause, the preempted alternative is not—and any analysis that does not yield that answer is in bad trouble" (Lewis 1987, p. 36). Nevertheless, given the logic of the situation as we have described it, had c_1 not happened, it would not have prevented c_2 from happening, and c_2 would have caused e , so c_1 is not necessary for e in the circumstances. If causes are necessary on the occasion, c_1 cannot be the cause of e . "It is easy for me to say why the preempted alternative is not a cause: the effect does not depend on it. My problem is to say why the preempting cause is a cause, when the effect does not depend on it either" (*ibid*). Hence, the problem is *logically* analogous to the symmetrical case. They differ in our intuitions, and in how they have been dealt with in theories of causation. Several attempts to resolve this situation, in particular Lewis's, hinges on that there is a series of causal dependencies (perhaps called a process or causal chain) all the way from c_1 to e but not from c_2 to e . That there is no such causal chain from c_2 when c_2 doesn't happen is obvious. But it's not obvious, at least to me, that there is *always* an intermediate causal link between any cause and its effect. If the assumption is that there is one just when we need it—that is to say, in cases of preemption—then this appears severely ad hoc.

Either way, cases of so-called late preemption introduce new difficulties. An often used example has two persons, say Kim and Jim, each throwing a rock at the same window, but Kim throwing ever so slightly before Jim. Kim's rock breaks the window, and Jim's passes through the resulting hole. Had Kim not thrown her rock, Jim's rock would have broken the window. Both throws happened, and Kim's throw is not necessary for the breaking of the windows in these circumstances. Here there is no intermediate event in a causal chains from Kim's throw to the breaking of the window, that happens at a time when the causal chain from

Jim's throw has already been preempted, but before the effect takes place. Nothing interrupts the causal chain from Jim's throw, only the actual breaking of the window, the effect itself, prevents Jim's rock from breaking it. Lewis considers an appeal to fragility for these cases (Lewis 1987, p. 40). The window breaks ever so slightly later than it would have, had Kim not thrown her rock. But he ultimately rejects this solution, since he can imagine situations where we must make the effect so fragile that we include some subtle feature of it that depends on the preempted cause, and redundancy is not escaped after all. It is beyond the scope of this text to recount the many important proposed solutions to late preemption, or even just Lewis's solution (but see Paul 2003 for a good summary). Since the approach to cases of apparent redundancy presented here does not distinguish substantially in its treatment of early and late preemption, nor between other types of redundancy, such as preempted prevention and trumping, we should get to it now.

6. TWO STRATEGIES FOR REDUNDANT CAUSATION

6.1. **Fact causes: a recap.** To recapitulate: I have argued that facts are plausible causal relata, based on that there is no consistency in what kind of things we in ordinary causal talk designate as the cause on some occasion, that facts can stand in the right nomic and logical relations to be causes and effects, and deal with absence causes better than events do. I have also suggested that they are metaphysically plausible when viewed as properties of worlds. (This is without question the most contentious part of my proposal.) These can in turn be modeled as sets of worlds, which leads to a picture of facts, whether presence- or absence-, as multiply realizable. We saw that writers on mental causation have both recognized the role of multiply realizable causes, and generalized these beyond the domain of mental causation. I've tried to introduce a further generalization, to multiply realizable facts, of which mental facts would be a special case, and connected these to

a predictable notion of relative specificity between sentences expressing facts, that follows logical strength and entailment.

Another thing these facts can do that events cannot is allow for disjunctive causes (and effects). We have already seen what appears to be an uncontroversial example. Some causes are multiply realizable. Correspondingly, some facts are naturally partitioned such that each part is a distinct way for the fact to be realized (such as a distinct time interval). For every partition, some part contains the actual world—that is, some part is the actual realizer of the fact in question. We then suggested that when both cause and effect are multiply realizable, the realizer of the cause, however chosen, will cause a realizer of the effect, of corresponding specificity. But not just corresponding in specificity—also in the aspect of realization on which the partition was made. The second-precision timed realizer of the stone throw causes the second-precision timed realizer of the window breakage. More intelligible: “the window broke because a stone was thrown at it,” but “the window broke at t because a stone was thrown at it at $t - \epsilon$ ” (for t of some precision).

That is to say, the question of more or less specific realizers of a fact is not like more or less modally fragile events. In expressing facts we have at our disposal the full resources of the language.

6.2. Preemption and finding the right effect. Kim throws her rock at the window. Jim doesn't bother to throw his rock, but would have thrown it at the window if Kim hadn't thrown hers. It seems as though Kim's throw isn't necessary under the circumstances for the breaking of the window—nevertheless, I wholeheartedly agree with the intuition that Kim's rock caused the window to break. Why? I think it is because Kim's throw caused the realizer of the window's breaking, and Jim's throw, since it didn't even happen, didn't. This is a matter of recognizing that the real effect of Kim's throw was something more specific than the window breaking

in general, but this intuition is already evident in past attempts to take the effect of the intuitive cause to be an event more fragile than first supposed. It seems as though that strategy works for us—and does so partly because facts, as opposed to events, are not indiscriminately more or less specific.

What if Jim *had* thrown his rock, just a little later than Kim threw hers, as in the case of late preemption? Then both rock throws happen, but still only one of them caused a realizer of the breaking of the window, and that was Kim's. That is, Kim's throw, and not Jim's, was necessary and sufficient for a realizer of the breaking of the window. Take the realizer in question (there will be many candidates) of "the window broke" to be "the window broke at t ," with the precision of t sufficiently high. If Jim's rock throw had broken the window after t , then it was not sufficient for that realizer.

The general strategy for cases of preemption is then to find the fact that the intuitive causal fact *alone* was sufficient and necessary for, and that is a realizer of the supposed effect. If we are successful, we can take that to validate our causal intuitions for that example.

6.3. Symmetric overdetermination and interpreting disjunctive causes. What if there is no such realizer of the effect, that one of the redundant alternatives is uniquely sufficient and necessary for? These would be strictly symmetrical cases of overdetermination. Lewis suggests that it is very unlikely that there are such in the actual world, and I think so too. In the case of the firing squad, the actual death from nine bullets (given that all the soldiers fired) constitutes a different fact from the merely possible eight-bullet death. The nine-bullet death is then caused by the *conjunction* of the actual firings of each soldier, and there is no redundancy. But if the causal concept makes sense also in fictional worlds, it may be enough that there is such a world, perhaps one where magic works, in which we can intelligibly

talk of causes, and where there are such genuinely and strictly symmetric causations. (They may then also contain cases where there are no intermediate links in a causal chain from cause to effect.) In such situations we cannot regiment our effect sentence to find the uniquely sufficient cause, but must instead try to understand the disjunctive cause sentence directly. Firstly, we saw in our definitions of facts and specificity that disjunctive facts are not special. Almost every sentence is equivalent to a disjunction of more specific sentences. So we might just take the disjunction as expressing the cause fact, see this as a result, and leave it at that. But this is not completely satisfying—surely there is some way of making sense of at least some of these disjunctions? There is, in many cases, if not all. We said that a disjunction can be taken to express a less specific fact than its disjuncts, under certain conditions. Sometimes there is also a non-disjunctive sentence that is equivalent to the disjunction in question, and intuitively less specific than the disjuncts. This was straightforwardly true when adjusting specificity was a matter of adding or removing predicates, adverbs, and adverbial phrases, but it is also true when we are dealing with disjunctions. If two magicians *R* and *S* both cast a spell that will turn the prince into a frog at midnight (this example is a simplification of one by Jonathan Schaffer, where he uses it for other things), then we can stipulate that there is no difference at all in the effect as it would happen if only one, or only the other, or both of the magicians cast the spell. This is a situation where we would have to accept that the disjunction, “*R* cast the spell or *S* cast the spell” is the cause. But this disjunction is equivalent to an existential generalization, restricted by context: “Some magician [of *R* and *S*] cast the spell,” that we may well find intuitively more acceptable than the disjunctive form. There is such a cause also for the breaking of the window and the execution of the prisoner, when these effects are taken sufficiently generally. The fact that some person [of Kim and Jim] threw a stone caused the window to break. And in the early preemption example, Kim’s throw was the realizer of that more general cause.

Generally, then, if none of the redundant alternatives is such that it uniquely was sufficient for a realizer of the effect, we must say that the cause was the disjunction of the redundant alternatives that obtained on the occasion, and try to achieve some intuitive understanding of this disjunction as a cause—but we are equipped to do so in many cases.

6.4. Multiple realizability again. I'll give one more example of the strategy of taking a disjunction of alternatives as the cause of some effect. The issue for Yablo in the quoted passages above was a competition in causal relevance between a multiply realizable event and its realizer. Consider again Sophie, pecking, as she was taught, at a red triangle that happens to be scarlet. By the exclusion principle, we might say that the causal sufficiency of the realizer, the fact that the triangle is scarlet, excludes that the redness was the cause of Sophie's pecking, contrary to how the story of her training was told to us. Another alternative was to invite causal redundancy. It's hard to believe, perhaps, that there would be a change in the specifics of Sophie's pecking act depending on what shade of red the triangle is, so let's assume that regimenting the effect is not an option. What remains is to treat the disjunction "The triangle is red or the triangle is scarlet" as the cause of Sophie's pecking at the triangle. But in this case, the disjuncts are not independent. In fact, the latter entails the former, even if not logically then just in virtue of that the fact expressed by the latter sentence *is* a realizer of the fact expressed by the former. (The worlds in which the triangle is scarlet constitute a subset of the worlds in which it is red.) This means that the disjunction is of the general form $\phi \vee (\phi \wedge \psi)$ —where ϕ corresponds to the less specific sentence "the triangle is red"—which is equivalent to ϕ . Generally, any disjunction is equivalent to its logically weakest, that is to say least specific, disjunct—which may be one or more disjuncts—so the strategy of taking the redundant alternatives will give us the least specific sufficient alternative as the cause.

7. ISSUES: STRICTURE, METAPHYSICS, AND LOGICALLY EQUIVALENT SENTENCES

This theory demands a lot of our causal claims. In fact, if our everyday causal statements are taken literally, it is unlikely that we ever make any true causal claims. This could certainly be seen as a problem. The theory (as far as it goes) does not validate everyday causal talk, when that is taken verbatim. On the other hand, we *know* that what we say in everyday contexts is imprecise and rarely exactly true. (But not all false claims are equal.) We know this, because we recognize contexts in which we must better ourselves, such as when giving a properly scientific explanation. Moreover, with sufficient theoretical demands on stricture, we may detect patterns in how we talk loosely about causes, that could be explained. It seems to me, for example, that we have a tendency to tolerate in particular underspecified effects and overspecified causes. This might be interesting.

As I have already suggested, the metaphysics of the causal relata presented here is, to say the least, contrary to current mainstream views. The demand that causes “exist in the world” may or may not be seen to be met on this account. However, the history of science and of our theoretical pursuits in general seem to show that the intuitiveness of concepts employed weigh lightly compared to the expressive power and pragmatic worth of theories. (Consider cases of space, time, and action-at-a-distance. Perhaps also calculus.) Partly for this reason, and partly because I have no better suggestion for understanding causes while achieving the theoretical power that appears to me required to reason fruitfully about causes, I tend to focus on whether the theory gets truth values of causal claims right rather than on metaphysical intuitiveness (without completely ignoring the latter—we continue to have beliefs about space, time, and action-at-a-distance even if they are revised). I expect significant resistance on this front, and look forward to criticisms and suggestions.

On a related note, in Laurie Paul's (Paul 2004), in which she presents a theory of aspect causation which appears to me a significant improvement compared to the most influential event-relational views, the tension in fine-grained event theories (Kim's and Lewis's) between on the one hand the need for enough theoretical power to be able to access what is causally relevant in a singular causation, and on the other hand the need to live up to metaphysical intuitions about causal relata, becomes exceptionally clear. This is largely because Paul puts pressure on the parts of these theories where they have problems delivering, and hers can. In taking aspects of events to be causal relata, instead of including event properties in the theory but nevertheless take causation to relate events as a whole, she continues, in my mind, an approach toward a fact-relational view begun by Kim and Lewis. (No doubt they would disagree.) I take the final step—making causal relata a full Boolean algebra mirroring the entailment relation between sentences—at the expense of some metaphysical intuitiveness. (This makes it different to Mellor's theory, which does not view causation as a relation at all.) Having this structure also makes possible a uniform treatment of absence and disjunctive causes.

A final concern. If what matters to the truth of a causal claim is that the antecedent sentence expresses the right fact, then we cannot distinguish between logically equivalent sentences. This means that we have no condition for excluding causally irrelevant things from being *mentioned* in the antecedent. For, if ϕ expresses the cause, then so, as we have already observed, does $\phi \vee (\phi \wedge \psi)$, and ψ can be anything—for example the conjunction of every sentence in the Bible. This is because we do not want to try to introduce *syntactic* conditions—that tends to not end well (see Kim's criticism of Mackie, in (Kim 1971)). What we can do is easily imagine that such conditions are introduced on the pragmatic level—that of logically equal statements, a more efficient one will be preferred when possible.

REFERENCES

- Jonathan Bennett. *Events and their names*. Hackett, 1988.
- Donald Davidson. The logical form of action sentences. In Nicholas Rescher, editor, *The Logic of Decision and Action*. University of Pittsburgh Press, 1967a.
- Donald Davidson. Causal relations. *Journal of Philosophy*, LXIV(21):691–702, November 1967b.
- John Gibbons. Mental causation without downward causation. *Philosophical Review*, 115(1):79–103, January 2006.
- Jaegwon Kim. Causes and events: Mackie on causation. *Journal of Philosophy*, 68:426–41, 1971.
- Jaegwon Kim. *Supervenience and Mind: Selected Philosophical Essays*. Cambridge University Press, 1993.
- David Lewis. Causation. In *Philosophical Papers Volume II*. Oxford University Press, 1987.
- Kirk Ludwig. Adverbs of action and logical form. In O'Connor & Sandis, editor, *A Companion to the Philosophy of Action*. Blackwell Publishing Ltd., 2010.
- J. L. Mackie. Causes and conditionals. *American Philosophical Quarterly*, 2.4: 245–55, 261–4, October 1965.
- D. Hugh Mellor. *The Facts of Causation*. Routledge, 1995.
- L. A. Paul. Aspect causation. In Hall & Paul Collins, editor, *Causes and Counterfactuals*. MIT, 2004.
- Ned Hall & L. A. Paul. Causation and preemption. In Clark & Hawley, editor, *Philosophy of Science Today*. Oxford University Press, 2003.
- Greg Restall. One way to face facts. *The Philosophical Quarterly*, 54(216):420–426, July 2004.
- Jonathan Schaffer. Causes need not be physically connected to their effects: The case for negative causation. In Christopher Hitchcock, editor, *Contemporary*

Debates in Philosophy of Science. Blackwell Publishing Ltd., 2004.

Stephen Yablo. Mental causation. *The Philosophical Review*, 101(2):245–278,
April 1992.