The Varieties of Relevance

Bradford Skow

Abstract. Causal relevance comes in many varieties, and causation is only one of them. There is direct and indirect relevance, and positive and negative relevance, but only the directly and positively relevant things are causes. Relevant non-causes include background conditions, enablers, ennoblers, and preventers. By distinguishing enabling from causing we can defang switching examples, so that they no longer threaten the transitivity of causation. And by distinguishing preventing from causing a negative (and I have a lot of arguments that they differ) we can clear a path to denying that double prevention is causation. But I don't issue a blanket denial: double prevention sometimes is and sometimes isn't causation, and I describe when.

1 The great misclassification

Non-causes are often misclassified as causes, and everyone knows it. In one famous example, economists observed the strong (negative) correlation between inflation and unemployment, and advised governments that increasing the inflation rate could cause unemployment to go down. This mistake produced the high-inflation high-unemployment 1970s (see Landsburg 2012, 272). But correlations are not the only kind of non-cause mistaken for a cause. Also mistaken for causes are things that are close to being causes; things that are, in some way, causally relevant to the outcome. My aim is to describe and categorize the varieties of causal relevance, compare them to the varieties of relevance in some other domains, and explain why recognizing them is important. The benefits will include a defense of the transitivity of causation, and a diagnosis of when double prevention is, and is not, causation. My view might be called pluralism about causal relevance, and this name brings to mind pluralism about causation: the thesis that causation itself comes in many varieties. I'll say something about the relation between these in section 12.

Exhibit A in my case for causal relevance without causation is the presence of oxygen in a room, when a match is lit. Many philosophers classify this, if only implicitly, as a cause of the match's lighting, but it is not. Causes have to be events, and the presence of oxygen is not an event. An event happens only when, and in virtue of the fact that, something does something. When I strike a match I thereby do something (what I do is strike a match), so in virtue of my doing it there is an event, the striking. But when I am human I do not thereby do anything (it is not true that what I do is "be human"), so even though I am human no event is happening in virtue of this fact. Applying this criterion to our case, no event happens in virtue of the oxygen's being present in the room, since *being present in the room* is not something the oxygen is doing. So the presence of oxygen isn't even the right sort of thing to be a cause.[1]

This argument relies on heavy-weight premises about what can be causes, and about the conditions under which events happen, which I have defended elsewhere, and which I will say a more about below.[2] Another argument leverages the distinction between being directly and being merely indirectly relevant to why something happened. Causes, and only causes, are directly relevant. The striking of the match is directly relevant to why the match lit. The oxygen, or the presence of oxygen, on the other hand, is only indirectly relevant. It is only relevant to why the match lit because it is a condition in which the striking of the match can cause the lighting. Without the oxygen, the striking could not have that effect. So the oxygen bears on why the match lit by bearing on the relation between the lighting and one of its causes; and that is an indirect way of bearing on why the match lit. Since it is only indirectly relevant, the presence of oxygen is not a

[1]Philosophers who implicitly mis-classify the presence of oxygen as a cause include David Lewis, whose theory of events in (Lewis 1986b) entails that the presence of oxygen is an event, and whose theory of causation in (Lewis 1986a) entails that it is a cause of the lighting. Ned Hall never articulates a theory of events, but he frequently treats things like a forest's being present, or a bit of train track's being present, as events and as causes (2004, 231; 2000, 207). The structural equations approach to the metaphysics of causation, which has many advocates, but see (Woodward 2003) for one canonical statement, says that (token) causes are cases of a variable taking on a value, and allows a model to contain a variable that takes value 1 when oxygen is in the room and 0 when it is not.

[2]See chapter 2 of (Skow 2018), where I also defends a version of the argument to follow. The theory of events is not original with me. It has its roots in Davidson's use of events in the semantics of action verbs (1967), and a similar theory receives a thorough defense in (Steward 1997).

cause. The presence of oxygen exhibits a variety, or mode, of relevance to the match's lighting, and to why the match lit, different from the one exhibited by causes of the lighting.

2 More on direct vs indirect relevance

My second argument that the presence of oxygen is not a cause distinguished direct from indirect causal relevance. One might object that this distinction is not real. But a distinction between direct and indirect relevance is all over topics outside of causation, and if it exists outside causation surely it also exists inside.

One place it exists is in inquiry into normative reasons. The canonical articulation of such a reason takes the form "that P is a reason for X to Z": that Jones is in need is a reason for me to help her; that the store down the street sells chocolate is a reason for me to go there. Now a normative reason answers a certain kind of why-question. A reason for someone to do something is part of the answer to why they should do it. If that P is a reason for X to Z, then that P is part of why X should Z. Part of why I should go to the store is that it sell chocolate. The reasons there are for me to go to the store, then, are directly relevant to why I should go.

By contrast, other facts, in particular facts about what I want, like the fact that I want to eat some chocolate, are (in the usual case) only indirectly relevant. My wanting to eat chocolate is relevant to why I should go only because it makes it the case that the fact that the store sells chocolate is a reason for me to go. More generally, someone's desires are (weird cases aside) indirectly relevant to why they should do this or that; they are relevant only by helping determine whether some fact is a reason for them to do this or that, and so whether that fact is directly relevant to why they should.[3]

[3]Note that I am not making the strong claim that whenever something is a reason for someone to act, some desire they have makes this the case. This claim or something like it is sometimes called "the humean theory of reasons"—see for example (Schroeder 2007).

Schroeder also argues strenuously for something like the distinction between direct and indirect relevance: he holds that the best version of the humean theory says that, if X's having a desire D makes it the case that R is a reason for X to Z, it does not automatically follow that X's having D is also a reason for X to Z.

The appearance here of the distinction between direct and indirect relevance is plausible on its face. The distinction is also built in to one popular theory of reasons to act. The theory says, roughly speaking, that the fact that P is a reason for X to Z if and only if a belief that P could be a premise in good practical reasoning that moved X to Z.[4] Maybe X does not in fact believe that P—maybe I've been told the store doesn't sell chocolate. Still, if I have some mental states M (including, say, a desire for chocolate), and if it would be good practical reasoning for me to be moved to go to the store by M-together-with-the belief that the store sells chocolate, then that the store sells chocolate is a reason for me to go. On this theory good practical reasoning may and often does involve desires helping (as parts of M) to move one to act. But in good practical reasoning beliefs about one's desires need not help move one to act, and in fact this is the normal case.[5] So, again, the fact that I want some chocolate is not itself a reason for me to go to the store. But still—that I want some chocolate is relevant to why I should go! It is indirectly relevant.

Indirect relevance is also there in epistemology, in the distinction between rebutting and undercutting defeaters. I might get some evidence E, which supports an hypothesis H, and then get some more evidence D that defeats my support for H. Defeat comes in direct and indirect varieties. The new evidence D might directly support an alternative to H, and support it strongly enough that overall my evidence is against H: then D is a rebutting defeater. Jones tells me that there is a rabbit in the next room; that's evidence that there is. I then go in and see for myself, and don't see a rabbit, or any of the paraphernalia that usually accompany them. This new evidence supports the no rabbit

Jonathan Dancy, as part of his defense of particularism about reasons, also argues for a distinction between reasons to act, and factors that determine whether something is a reason, or determine the strength or valence of a reason (Dancy 2004). For him, facts about what the agent desires are not the only such factors.

The weird cases I have set aside are ones like this: a philosophy department promises a high salary to applicants who want to teach large introductory classes. In this case, the fact that I want to teach large introductory classes is a reason for me to apply to the job. It is directly relevant to why I should apply. (That I want a higher salary is, though, only indirectly relevant: it makes the fact that I want to teach large classes directly relevant.)

[4]Among many others, (Setiya 2014) defense a theory like this.

[5]This is a thesis about good practical reasoning; what about the practical reasoning we actually engage in, which may not be good? Pettit and Smith (1990) argue that here also beliefs about one's desire rarely help move one to act.

hypothesis more strongly than Jones's testimony supported the rabbit hypothesis, and in virtue of this the new evidence defeats the evidential support Jones' testimony gave to the hypothesis that there was a rabbit.

Undercutting defeaters, by contrast, do not directly support an alternative to H. What they do, in the first instance, is undermine my original evidence E's support for H. If, when facing a wall with my eyes open and inquiring into its color, all I have to go on is how the wall looks, then that the wall looks red to me is evidence that it is red. But if I learn that the wall is illuminated by a red light, this new fact defeats my original evidence's support for the red hypothesis. But it does not do this by directly supporting some other hypothesis about the color of the wall. It does it by directly destroying the support my first evidence was giving the red hypothesis.

These differences in what the kinds of defeater bear on most directly translates into a difference in their mode of relevance. Rebutting defeaters are directly relevant to whether I should believe H. Undercutting defeaters are (only) indirectly relevant. What I see when I enter the allegedly rabbit-inhabited room is directly relevant to whether I should believe that there is a rabbit, because it directly supports the "no" answer. But the fact that the wall is illuminated by red light is only indirectly relevant to whether I should believe the wall is red: it is relevant only by making the fact that the wall looks red irrelevant.[6]

In these other manifestations of the direct/indirect relevance distinction there is a special role reserved for things that are directly relevant. The things directly relevant to why I should Z are the reasons for me to Z; the indirectly relevant things, like my desires, are not reasons. The things directly relevant to whether I should believe H constitute my evidence for and against H. The indirectly relevant things, like undercutting defeaters, are not evidence for or against H. What they are is evidence that something I earlier took to be evidence for H in fact is not. It is the same with causation: the things directly relevant to why E happened are the causes of E; the indirectly

[6] A canonical source for the distinction between rebutting and undercutting defeaters is (Pollock and Cruz 1999). Undercutting defeat is not just a kind of indirect relevance, it is also a kind of negative relevance: it prevents some evidence from supporting an hypothesis that it otherwise would. I will have more to say about negative relevance later, when I discuss prevention.

relevant things are not causes, though they are, as we might say, "relevant to the causing."

Before going on, a note: *being directly relevant to E,* the notion I've been explaining, is distinct from *being a direct cause of E,* and indeed indirect causes, being causes, are directly relevant to their effects. Indirect causes act by way of intermediate causes, as when the lighting of the fuse causes an explosion by way of causing the gunpowder to ignite. This is quite different from indirect relevance.

3 Background conditions as indirect mode

The presence of oxygen is a non-cause of the lighting of the match that is still, indirectly, relevant to why the match lit. But, I will argue, there are many kinds of indirect causal relevance. A good name for the kind that the presence of oxygen exhibits is *background condition.*

The standard view about background conditions is that the notion is, in Hall's words, "greatly infected with pragmatics":

> When delineating the causes of some given event, we typically make what are, from the present perspective, invidious distinctions, ignoring perfectly good causes because they are not sufficiently salient. We say that the *lightning bolt* caused the forest fire, failing to mention the contribution of the oxygen in the air, or the presence of a sufficient quantity of flammable material. But in the egalitarian sense of "cause," a complete inventory of the fire's causes must include the presence of oxygen and of dry wood. (2004, 228)

I think the opposite is true: there are good metaphysical (and so non-pragmatic) reasons for excluding the oxygen from the list of causes. Summarizing the differences that drove my two arguments, a background condition to C's causing E is (i) a state (rather than an event), that is (ii) a reason why C caused E. Clause (ii) captures the sense in which background conditions are indirectly relevant.[7]

[7]I defend this theory, and argue against alternatives, in (Skow 2018). A useful table summarizing alternative ways of drawing the cause/background condition contrast is in (Goldvarg and Johnson-Laird 2001). (Their table is described as summarizing ways to draw the cause/"enabling"

But more should be said about the difference between states and events.

4 Event/state and activity/stative

Here again is my existence condition for events:

> Events: if something did something, then in virtue of that fact an event occurred, and if an event occurred, that's so in virtue of the fact that something did something.[8]

If a verb V can be put into "What he/she/it did was V" to produce something grammatical, then that verb denotes[9] something that can be done—acts, we might call them. I will call these verbs activity verbs, and verbs that are not activity verbs, stative verbs.[10] States go with stative verbs in the way events go with activity verbs. The existence condition for states is harder to frame than the one for events; one way is this:

> States: if something Vs, where "V" is a stative verb, then in virtue of that fact a state obtains, and if a state obtains, that is so in virtue of the fact that something Vs.[11]

condition contrast, and while I will later distinguish background conditions from enabling events, they are not distinguished by these authors.)

[8]The quantifiers take wide scope: for any X and Y, if X did Y, then in virtue of the fact that X did Y, an event occurred. The case where X is plural is also allowed. I am using "event" broadly; it also includes processes. While standing is something that may be done, some linguists say that "Jones stood" (and sentences with other, similar "doing" verbs) entails the existence, not of an event, but of a state, just a special kind of state: a "dynamic state" (Bach 1986, 6) or a "Davidsonian state" (Maienborn 2008, 109). Verbs that do not report doings, like "be tall," they say entail a different kind of state, a "static" or "Kimian" state. But I use "event" to include dynamic/Davidsonian states; the only states are "static/Kimian" states.

[9]In the sense of "has as its semantic value"—verbs of course are not referring expressions like "Barack Obama" or "him."

[10]Some linguists use "activity verb" in a more restricted sense, so that only some non-stative verbs are activity verbs. To keep things simple, I am using "activity verb" and "non-stative verb" interchangeably. Important early papers distinguishing stative from (a variety of) non-stative verbs are (Vendler 1957) and (Kenny 1963). (Dowdy 1979) contains a then-comprehensive list of tests for the distinction.

[11]Again the quantifier takes wide scope and the plural case is allowed. Should we recognize states even for the narrow-scope reading—a state that goes with [someone is tall], separate from the one that goes with [Jones is tall]? Maybe. Nothing in this essay will turn on this question. A

A problem with this statement is that English does not have a stative verb for every possible stative verb meaning, so it is insufficiently general. If we had a variable V that could occupy verb position, and moreover ranged only over things that can be denoted by stative verbs, then we could write the condition like this:

> States: if $\exists V\, \exists x\, (x\, V)$, then in virtue of that fact a state obtains, and if a state obtains, then $\exists V\, \exists x$ (it obtains in virtue of the fact that $x\, V$).[12]

Now: the verb in "Oxygen is present in the room" is stative, so the state corresponding to it meets the first condition on being a background condition. As for the second: that that state obtained is indeed a reason why the striking of the match caused it to light. Without the oxygen, the striking would not have caused the lighting.

The distinction between activity verbs and stative verbs runs deep in natural languages. This distinction, and the notion of doing something that it relies on, are real and metaphysically important. If the distinction is less familiar to philosophers than to linguists, it is because philosophers learn too early to formalize their arguments in artificial languages where the distinction is invisible. In those languages a predicate has some number of argument places, but other than that all predicates are grammatically indiscernible. They are not divided into activity predicates and stative predicates. These artificial languages are not eliminating an irrelevant distinction; they are blinding themselves to a relevant one.

Some philosophers accept something close to the thesis that an event happens whenever something satisfies a predicate, any predicate, and so in effect treat the distinction between activity and stative verbs (and predicates) as superficial and of no metaphysical significance, akin, maybe, to the distinction between predicates of English and predicates of Japanese, or the distinction between verbs that begin with "s" and those that do not.[13] The latter two distinctions are obviously

referee asked whether a vacuum is a state, on this view. If a vacuum is a region of space that is empty, the answer is yes; "is empty" is stative.

[12] I suspect that I don't actually need to "reify" states; the arguments and conclusions I want to reach could probably be reached by always talking about stative verbs. But it is easier this way. (Maienborn 2008 argues that the semantics of stative verbs requires reifying states.)

[13] Kim (1993, 33) is relatively explicit about this. Lewis's theory (Lewis 1986b) also has this

metaphysically (and grammatically) irrelevant, but the first is not.

I've given one reason already: if V is an activity verb phrase, then "X Ved" is a possible answer to "What did X do?," and so denotes something that can be done. If V is a stative verb phrase, "X Ved" makes no sense as an answer to this question, and does not denote something that can be done. And the notion of doing something is metaphysically important.

Another thing that makes the activity/stative distinction metaphysically important is its connection to locations in space. Activity verbs accept locative modifiers with spatial interpretations, while stative verbs do not. For example, "The bomb exploded in the courtyard" is interpretable when the modifier has a spatial reading, but "Jones was 6 feet tall in the courtyard" is not. The immediate metaphysical consequence is that events, which go with activity verbs, have locations—every event happens somewhere—while states, which go with stative verbs, do not. I earlier asserted that only events can be causes, and you might have wanted me to say more about why. There are many reasons, but here is one: states cannot be causes because causes and effects are things in space and time, and states are at best only in time.[14]

The second, stative, example in this argument ("Jones was 6 feet tall in the courtyard") is admittedly tricky, because the sentence is interpretable—but only when "in the courtyard" is read, not as a locative modifier, but as a temporal modifier that restricts the past tense, making it

consequence. Neither really says that the happening of an event goes with the satisfaction of just any predicate: Lewis for example restricts the thesis to "relatively natural" predicates, in the sense of natural articulated in (Lewis 1983), but naturalness is orthogonal to the activity/stative distinction.

[14]The inference that states lack spatial locations is strongest if it is assumed that stative verbs have hidden argument places for states, in the way that (neo-)Davidsonians assume that activity verbs have hidden argument places for events. For then at the level of semantic interpretation there is a state variable available for the location predicate to attach to, and so the ungrammaticality of adding locative modifiers must be due to states' inability to have locations (see Maienborn 2005, 2008). But I think that the argument has force even without these semantic assumptions, on which I take no stand.

Another argument that only events can be causes is that in the form "X caused Y to Z by Ving"—"Jones caused the window to break by kicking it" is an example of an instance—only activity verbs (and so only verbs that go with events) may go in for "Z" and "Ving"; see (Skow 2018) for more on this argument.

equivalent to "When Jones was in the courtyard, he was 6 feet tall."[15] So when does a locative modifier "in W" have a spatial interpretation in a sentence "S in W," and when does it have some other function, like restricting a quantifier? One useful criterion: when the occurrence of "in W" can be questioned with "where." You can question "The bomb exploded in the courtyard" with "Where did the bomb explode?"; but you cannot question "Jones was 6 feet tall in the courtyard" with "Where was Jones 6 feet tall?"[16]

Related to this is a third metaphysically-relevant difference between activity and stative verbs: only activity verbs can appear in complements to perception verbs (Maienborn 2005). "Smith saw the bomb explode" is grammatical but "Smith saw Jones be 6 feet tall" is not. This is to be expected if only events have spatial locations: you can only perceive what happens in space.

5 Ennoblers

States, and the event/state contrast, came up when I said that background conditions are states, and so are never causes, even though they are often misclassified as such. The air's being still is not a cause, but is instead a background condition, when I draw a bow and release the arrow at a certain angle, and thereby cause the arrow to hit the bullseye. The valve on my radiator's being open is not a cause, but is instead a background condition, when I turn up my thermostat and thereby cause my radiator to get hot. The existence of background conditions shows that being a cause of E is not the only way to be causally relevant to the occurrence of E. What's more, *cause* and *background condition* are just two of a whole menagerie of kinds of causal relevance. If you want to cultivate a sensitivity to the varieties of causal relevance, a good place to start is Stephen Yablo's paper "Advertisement for a Sketch of an Outline of a Prototheory of Causation" (2010).

[15] If this still sounds funny, it is because we know that people move around much faster than their heights change. But this shows only that the sentence is odd to assert (since it may carry the false implication that Jones was some other height before he was in the courtyard), not that it is uninterpretable.

[16] See (Maienborn 2001, 6). It is, of course, intelligible to ask, "Where was Jones when he was 6 feet tall?"; but this questions, not the target sentence, but "Jones was 6 feet tall when he was in the courtyard," where the modifier restricts a quantifier (a tense).

Yablo begins with Plato, or maybe it was Socrates, who distinguished between *the cause* and *that without which the cause would not be a cause,* when he said to Cebes, in the *Phaedo,* "Imagine not being able to distinguish the real cause from that without which the cause would not be able to act as a cause" (99b). Yablo himself distinguishes between two distinctions this sentence could be used to draw. The first is "the distinction between causes and enabling conditions," where enabling conditions are "conditions that don't produce the effect themselves but create a context in which something else can do so; conditions in whose absence the something else would not have been effective" (98). Enabling conditions typically have the feature that "if you imagine them away, the cause...ceases to be *enough* for the effect" (98) Enabling conditions thus defined sound a little bit like background conditions; more on their relation in a minute. If enabling conditions make the cause "enough" for the effect, could there also be, Yablo asks, "conditions such that the cause ceases to be *required* for the effect, if you imagine them away" (98)? He answers yes, and calls them *ennoblers.* Yablo repurposes an example of Hartry Field's to illustrate ennobling:

> Billy puts a bomb under Suzy's chair; later, Suzy notices the bomb and flees the room; later still, Suzy has a medical checkup (it was already arranged) and receives from her doctor a glowing report. (98)

The bomb's presence is an ennobler, because "The bomb does not help Suzy's leaving to suffice for the glowing report [as an enabler would]; rather it makes Suzy's action important, required, indispensable" (99). He goes on: "What an ennobler contributes is just a raising of status. Suzy's removing herself from the room is elevated from something that just happens to something that *had* to happen, if Suzy was later going to be healthy" (99). In general, "An ennobler contributes by closing off potential routes to *e* [the effect], that is, all the routes not running through *c* [the cause]" (99). By raising the status of Suzy's leaving from something that merely preceded the glowing report to something that had to happen for the glowing report to happen, the bomb's presence ennobled the leaving into a cause of the report.

Interestingly, while Yablo first characterizes ennoblers in terms of what would happen in their absence (the effect wouldn't have required C, where C is an actual cause of E), he almost

immediately switches to, and then for the rest of the paper works with, a definition in terms of what would happen in their continued presence: he says that G ennobles C into a cause of E iff had C not happened and G still obtained, E would not have happened (this is later made more complicated). The bomb example fits both definitions: without the bomb, Suzy's leaving would not have been required for the glowing report (that's the definition in terms of absence); and if Suzy hadn't left, but the bomb still been there, she wouldn't have received a glowing report (presence). It's an interesting question whether these two characterizations of ennobling are equivalent; but that question is a rabbit hole that I will gladly side-step. The two characterizations certainly agree on the examples I will discuss.

What matters right now is that ennoblers are not causes. We know enough about them that this is obvious. "An ennobler," again, "contributes by closing off potential routes to *e*, that is, all the routes not running through *c*. This is anything *hurts e*'s chances" (99). It's hard to see how something that works against E's happening could be a cause of E. Ennoblers play for the other team.[17]

I admit to finding *ennobler* a trickier concept than *enabling condition*. Maybe helpful is looking at more examples. If I set up ten sling-shots to sling rocks at a window in quick succession at the push of a single button, and the first nine mechanisms jam (coincidentally) when I push it, then the jamming of those mechanisms ennobles the slinging of the tenth rock into a cause of the breaking of the window. Without the jammings, the slinging of the tenth rock was not causally necessary for the breaking; holding fixed the nine jammings, had the tenth sling not happened, the window would not have broken. The jammings are certainly not causes of the breaking—there's no pull to say that at all. But they are relevant to the causing of the breaking, because they help make something else into a cause. Examples with essentially this structure are easy to produce. We

[17]This phrase comes up a lot in discussions about ennoblers I've participated in, and I was sure that Yablo himself used this metaphor, but it turns out it does not appear in his writings. My research suggests that Daniel Muñoz came up with the phrase and mis-attributed it to Yablo. Still: it is something Yablo should have said, in the way that Rick in *Casablanca* should have said "Play it again, Sam," or Sally Field in her Oscar acceptance speech should have said "You like me, you really like me!"

have won free pizza for the evening, and it is on its way. I don't know this, so seeing nothing in the fridge I order pizza from a different restaurant. The first delivery guy is pulled over for speeding and the pizza he's carrying spoils in the back seat, never making it to my house. His detention ennobles my order into a cause of my family's eating; without the detention, my order would not have caused the eating, since we would have eaten the free pizza. Holding fixed the detention, had I not ordered, my family would have gone hungry.

So Yablo gives us *enabling condition* and *ennobler* to set alongside *cause* and *background condition*: four kinds of causal relevance. But do we really have four? As the list grows so does the danger that it is redundant. In fact it is: enabling conditions are background conditions under a different name. Enabling conditions, we are told, "don't produce the effect themselves but create a context in which something else can do so," and that's the defining feature of background conditions. The oxygen didn't produce the flame, the striking of the match produced the flame; but it was only in a context where oxygen was present that the striking could do that. So really we have a three-item list. And this shorter list is not redundant. We have already seen arguments that neither background conditions nor ennoblers are causes; and background conditions are distinct from ennoblers, since (in Yablo's terminology) background conditions make the cause enough, while ennoblers make the cause required.

6 Enablers

This shorter list is not redundant, but neither is it complete. Suppose I fill a previously-evacuated room with oxygen, just before you, standing in the room and holding your breath, strike a match, thereby lighting it. Your striking of the match caused it to light. The oxygen in the room was a background condition to the lighting: in virtue of its presence, your striking sufficed for the lighting. What about my filling of the room with oxygen? In doing that, I caused the onset of a background condition.[18] No doubt that is something causally relevant to the lighting; this kind of

[18]I didn't cause the background condition itself—the presence of oxygen—since this is a state and cannot be caused. But the onset of a state, or a state's beginning to obtain, is an event, and so

thing belongs on our list, a true fourth entry. But "cause of the onset of a background condition" is unwieldy; can we come up with a shorter, punchier name? We can, by retrofitting the term we just crossed off. Yablo's "enabling condition," a term for a kind of condition, suggests "enabling," a term for a kind of act. Enabling is what enablers do. In the ordinary sense, people who are enablers clear the way for someone else to indulge in their bad habits. An enabler will pay off your gambling debts so you can keep going to the track. An enabler will cover for you at work so you can keep hitting the bar at night without getting fired. If we generalize this idea, abstracting away from the requirement that enablers be people and enable other people, we get the idea of something that affects the conditions that determine the course of a causal process. By paying someone's debts the enabler ensures that the gambling does not lead to their being fired; it's the same with covering for a heavy drinker at work. But manipulating the conditions that determine what effects C has just is putting into place conditions that are background conditions to C's having its actual effects (where other conditions would have made C produce different effects). If that is what an enabler does, then enabling is the doing of it. Enabling is causing the onset of a background condition. In more detail, if something caused the onset of a background condition to C's causing E, then that thing enabled C to cause E, and so is an enabler (and conversely).[19]

I've gotten the term "enabler" from thinking about something Yablo wrote, but really Lawrence Brian Lombard had the term, and the concept, much earlier (Lombard 1990). He distinguished between "the causes of events" and "those conditions or states the obtaining of which merely makes it possible for one event to cause another" (201); those conditions or states are, in my terms, background conditions. And he then defined an enabler, more or less as I do, as "an event that causes a thing to be in a state that makes it possible for an event to cause *e* [the effect in question]" (202).

Lombard cited no precedents, but in fact use of the concept of an enabler goes back much

can be caused.

[19]Yablo actually uses the term "enabler" in his paper, and identifies enablers with both enabling conditions and with causes: "Enablers make a dynamic contribution. They help to bring the effect about" (99). That's not what background conditions do at all, nor is it what enablers, as I've defined them, do (my arguments for this second claim are yet to come). Yablo is another philosopher who mistakenly identifies some non-causes as causes.

farther. Philippa Foot in a 1967 paper identified enabling with "the removal of some obstacle which is, as it were, holding back a train of events" (1978, 26). If a stretch of track is blocked by some rubble and I remove the rubble, I on this definition enable the train's departure to cause it arrival at the station further on. But Foot's definition is narrower than mine, and I think the wider definition is better. Removing an obstacle to C's causing E is one way to enable C to cause E, but not the only way; creating a path to C's causing E is also a way to enable, and is a different way. If a ball is rolling down a mountain and I build a little ramp so that it jumps over the stream at the bottom and lands safely on the other side, then I say that I enabled the initial rolling to cause the landing, but I haven't removed some obstacle to the rolling's having that result. What would the obstacle be? The absence of a ramp? Did I, in building the ramp, remove the absence of a ramp? Is the sky full of obstacles to my death, full of "absences of meteors heading toward me," that thankfully no one has removed? No. Yet there is a kind of relevance to one thing's causing another that building the ramp and removing the rubble have in common.

So enabling belongs on the list of kinds of causal relevance. But are we sure that adding it doesn't make the list redundant? Well it is clearly different from ennobling. And certainly *causing the onset of a background condition* is distinct from *being a background condition*. For one thing, the first is something only events can do, while the second is something only states can be.

What about causation itself? Is enabling just a kind of causing? Well of course in a sense it is: by definition, if you enable C to cause E then you cause something (the onset of a background condition to C's causing E). But that's not what I mean. We are compiling a list of ways of being causally relevant *to E,* so the question that needs answering is whether enabling C to cause E is really a way to be a (another) cause of E. I am convinced that it is not. This claim might raise eyebrows. Background conditions aren't causes, I said, because they aren't events. But enablers are events, have to be events, since they cause background conditions to begin obtaining. So if an enabler of E isn't a cause of E, it isn't for the reason that background conditions aren't. What's more, enablers of E pass a standard counterfactual test for being causes of E. If you hadn't struck the match, it wouldn't have lit, and the striking is a cause of the lighting. Well similarly if I hadn't

filled the room with oxygen, the match wouldn't have lit; so shouldn't my filling of the room also be a cause?

But look closer and the similarities between enablers of E and causes of E are swamped by deep differences. For one thing, enablers fail to be directly relevant to the effect, the way causes are supposed to be. An enabler inherits its mode of relevance from the mode of the background condition it puts into place. And background conditions, again, are only indirectly relevant. So my turning of the knob is relevant to the lighting of the match only by helping put in a place a background condition (the presence of oxygen) that makes it possible for the striking cause the lighting. That's not the kind of direct relevance to the lighting that makes for causation.

Of course in special circumstances something that enables C to cause E might also be a cause of E. If I turn a knob to open a valve that lets oxygen fill the room, and my turning of the knob also creates a spark that lights a fuse connected to the match, and the fuse burns to the end, lighting the match, then my turning of the knob enables the spark to cause the match to light, and is also a cause of the lighting of the match, in virtue of being a cause of the spark. But here my turning of the knob, the enabler of the lighting, is not a cause in virtue of being an enabler. That's the thesis I assert: something that enables C to cause E is not, just by virtue of this fact alone, a cause of E.

Some evidence in the psychological literature supports the idea that enablers are indirectly relevant. In some experiments when subjects are told that something identified as an enabler happened, they don't expect the effect to have happened unless something identified as a cause also happened (see for example Golvarg and Johnson-Laird 2001, and Sloman et al 2009). One explanation of these results—the correct one, I think—is that people believe that an enabler is relevant to an effect's occurring only indirectly, by being relevant to a cause's power to produce it. (Of course as I use them here I intend "enable" and "enabler" to be quasi-technical terms, and the subjects in these experiments were not instructed to interpret them in my technical sense. But my technical meanings are not too far from these words' ordinary English meanings, so these experiments are still relevant.)

7 Depicting indirect relevance

Neuron diagrams are a popular device among certain metaphysicians for representing and thinking about causal relationships. They also perfectly exemplify the fact that background conditions and enablers are overlooked, ignored, and even suppressed in philosophical discussions of causation. In the simplest two-neuron neuron diagrams, shown in figure 1, neuron A is connected to neuron B by a stimulatory connection, meaning that if A fires, that will cause B to fire. That's what happens in the top diagram: blackening the circle representing a neuron represents that it fired. In the bottom diagram neither fires.

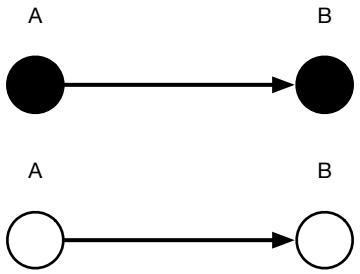

Figure 1

So the color of a neuron is used to represent the occurrence, or non-occurrence, of an event (a firing). How do the diagrams represent states, including background conditions? They don't. There is no standard way to represent the obtaining, or not, of a state. The metaphysician's standard tool for representing causal structures literally makes background conditions invisible. If we want a neuron diagram that illustrates a background condition to one thing's causing another, we're going to have to invent a scheme for representing that ourselves.[20]

Here's kind of a silly example. Suppose that neuron A's axon passes through a room on its way to neuron B. Whether the signal from A can make it through the room depends on what state

[20] A referee pointed out that the arrows between neurons represent states: the state that goes with "A and B are connected by a stimulatory connection." So some states are represented in these diagrams. But there is no standard way to represent these states changing, or to represent other causally-relevant states. Since metaphysicians use neuron diagrams to represent the causal structures of a variety of situations (not just actual neurons), this puts serious limits on what features of such situations the diagrams can represent, without resorting to ad-hoc devices. Any event can be represented by a neuron firing, but not just any state can be represented by an arrow.

obtains in the room: if it is filled with oxygen, the signal can pass through, otherwise it cannot. (So now we're departing even farther than usual from the way actual neurons work.) Figure 2 shows the setup: I'll use grayscale to indicate whether the room is filled with oxygen.

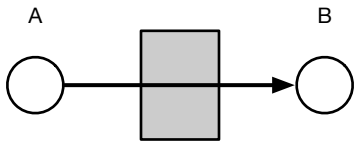

Figure 2

Now look at Figure 3: neuron A has fired, oxygen was present, so the firing caused neuron B to fire. While clear enough, the figure is misleading, since it uses the same thing—color (or black/white/gray shading)—to represent two different kinds of things: which events happen, and also what states obtain.

One way to separate these things out is to represent time as well as space. States typically persist through longish stretches of time, while events, or at least neuron-firings, happen very quickly. So each figure should contain a sequence of diagrams, read from top to bottom, that shows how things unfold over time. Momentary colorings represent events—firings—while persisting colors represent states. That's what we've got in figure 4: it shows that the state that is the presence of oxygen in the room obtained for a while, then neuron A fired, and that event took just a moment, and then the signal traveled down the axon leaving A and, because there was oxygen in the room, got all the way through the room, and caused neuron B to fire. (If persisting colors represent states, then the persisting whiteness of, for example, neuron B for the first part of the sequence should represent a state, right? Yes, indeed: it represents the state of the neuron's being dormant.)

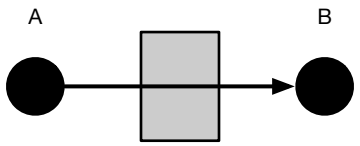

Figure 3

Now let's put an enabler into the scenario. Neuron C in figure 5 represents an oxygen pump,

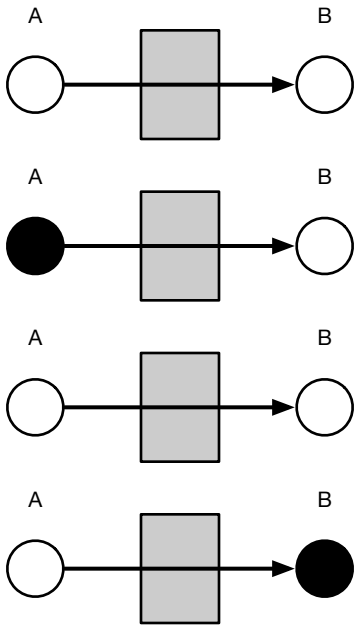

Figure 4

connected to the room by a hose, that can fill the room with oxygen. When neuron C is blackened, that represents that pumping is occurring. In figure 6 this happens: first C pumps oxygen down the hose, thereby causing the onset of a state (the presence of oxygen) in the room. After the pump shuts off and the room is full nothing happens for a moment. Then A fires, and because oxygen is present, it causes B to fire.

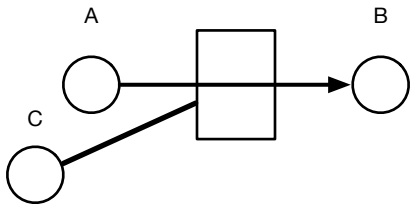

Figure 5

8 Enabling, causing, and locality

The scenario in figure 6 is the basis for another argument that enablers are not causes. The main premise is a locality requirement: a cause of E must be connected to E by a continuous sequence

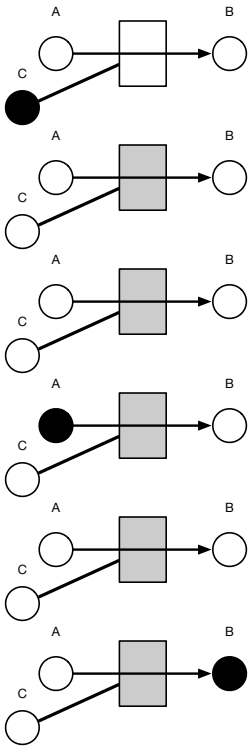

Figure 6

of intermediate causes (at least, this is so when we can ignore the kind of non-local causation that, some say, quantum mechanics permits). An enabler of E need not be; so enablers are not causes.

The locality requirement has a spatial and temporal part. The spatial part: if a cause C happens here, and its effect E happens way over there, then there is an intermediate event that is an effect of C, happens right next to C, and is itself a cause of E. This principle iterates, until we get an intermediate event that is an effect of C, is a cause of E, and happens right next to E.[21] But important for us is the temporal part: if C finishes happening at one time, and E begins happening some non-zero amount of time later, then there is an intermediate event that is an effect of C, happens right after C, and is itself a cause of E. (This principle also iterates.)

Now consider again the scenario in figure 6. The enabler, the pumping of the oxygen, causes the movement of the oxygen (an event), and the movement of the oxygen causes the onset of the state of oxygen's being present, and then, for a while, *nothing happens*. No events occur.

[21]This could be made more precise, but more precision isn't needed here. In iterating the principle I assume that causation is transitive; I discuss transitivity more below.

Eventually A fires, and a spatially and temporally continuous sequence of intermediate causes connects that firing to B's firing (these causes are events involving changes in the state of the axon connecting A to B and are not represented in the diagram), a sequence that passes through the room. But the onset of the state of oxygen's being present cannot cause any of these events, not even the ones that happen where it is, the ones that happen in the room, because there is a temporal gap between those events and the onset of the state, a gap that is not filled by any intermediate events.[22]

9 Indirect relevance and structural equations

Drawing neuron diagrams is, by now, an old way to conduct disputes about the metaphysics of causation. The new way is to present causal models using structural equations. Like neuron diagrams, these causal models also erase the distinction between causation and other modes of causal relevance.

A structural equation causal model of a situation consists of a set of variables and a set of equations relating those variables. In the simplest case, a variable represents an event: a variable X, for example, might take the value 1 if an explosion happens, and the value 0 when it does not. But variables need not be limited to two possible values: a variable might take the value e when an explosion releases e joules of energy, where e is any non-negative real number; and, as we will see, variables need not be limited to representing events.

The variables in a model come in two kinds, endogenous and exogenous. An exogenous variable's equation just sets that variable to a specific value: if V is exogenous, the equation says $V := v$ for a certain value v. An exogenous variable gets its value from "outside the system." An endogenous variable gets its value from inside: the equation for an endogenous variable Y deter-

[22]Lombard (1990, 201) also argues that enablers fail a temporal locality requirement on causation.

If you think the motions of the oxygen molecules constitute intermediate events, suppose that part of the scenario (especially the third diagram) happens at such a low temperature that the molecules are at rest. Being at rest is not something molecules can do, so being at rest generates no events.

mines its value from the values of other variables $X1, \ldots, Xn$. The equation $Y := f(X1, \ldots, Xn)$ is to be understood to mean that the Xi are all the direct causes of Y; this itself is taken to mean that the value of each Xi at least sometimes makes a difference to the value of f.[23] There are other constraints that a set of variables and equations must satisfy to constitute a possible causal model of a possible situation, but we will not need them.

To model a situation where I strike a match, causing it to light, let S be a variable with value 1 if I strike, 0 if I don't, and L be a variable with value 1 if the match lights, 0 if it doesn't. The equations of the model are

$$S := 1,$$

$$L := S.$$

Given a causal model, one wants to know what it takes for one variable X's taking on a value x to be a cause of another variable Y's taking on a value y. One popular answer says that $X = x$ is a cause of $Y = y$ in a given model M iff there is a sequence, or path, of variables $(X, V1, \ldots, Vn, Y)$ where $V1$ appears on the right-hand side of the equation for X, $V2$ appears on the right-hand side of the equation for $V1$, and so on, and, intuitively, there is a possible value for X with the property that, had X taken that value, while the variables off the given path from X to Y had been held fixed at their given values, then Y would have taken some other value.[24] In our simple example there are no off-path variables, so all it takes for $S = 1$ to be a cause of $L = 1$ is for it to be true that if S had

[23] $X1$ makes no difference to the value of f iff for all values $x2, \ldots, xn$ of $X2, \ldots, Xn$, $f(a, x2, \ldots, xn) = f(b, x2, \ldots, xn)$ for any possible values a, b of $X1$.

[24] This claim about what value Y would have taken if X had had value $x1$ is understood as a claim about what value Y does take in a specific alternative to M. First define a model $M1$ related to M like this: if the equation for X is not $X := x$, replace it with this equation. If W is a variable not in the given sequence, and $W = w$ in M, then the equation for W in $M1$ is to be $W := w$. So all variables not on the given path from X to Y are, in $M1$, held fixed at the values they have in M. Note that in $M1$ $Y = y$, just as in M. Finally, it is true in M that Y would have had a value other than y had X had a value other than x, and so true that $X = x$ is a cause of $Y = y$, when there is some alternative value $x1$ for X such that in the model $M2$ that differs from $M1$ only in that the equation for X is $X := x1$, Y takes on a value other than y. This theory of causation is given in (Hitchcock 2001a) and (Woodward 2003), among many others.

been 0, L would have been 0.

There are many philosophical questions about how these causal models are to be understood.[25] Here only one question is important: what sort of things can a variable in a causal model be used to represent? The standard answer, I take it, is there are almost no restrictions on what a variable can represent—a variable can correspond to any indirect question, from *whether event E happened* and *how much energy the explosion released* to *whether the room contained oxygen* and beyond.[26] Variables are given so much representational power because they have a big job to do. If you have something causally relevant that you want to represent in a causal model then the only way to represent it is to use a variable.

This makes it impossible to distinguish causes from background conditions, and from things with other modes of causal relevance, in these models. The oxygen in the room, when I strike a match, is, again, a background condition to the lighting, not a cause. But if you want to explicitly represent the presence of oxygen and its relation to the lighting in our causal model, you would have to do it by expanding the model to include a new binary variable O, with value 1 iff oxygen is present; the full list of equations would be

$$O := 1,$$

$$S := 1,$$

$$L := SO.$$

The variables O and S, for whether there is oxygen and whether the match is struck, enter these equations symmetrically. There is nothing in the model to tell you whether one is a cause or a

[25] One thorough treatment of such questions is in (Woodward 2003). Hall (2007) thinks these questions are harder than those who like to state their theories of causation using these models take them to be.

[26] "Beyond" might include variables like *what species this individual belongs to* and *how many people are in the room.* Some authors, like (Woodward 2003), formally allow all these variables, but deny that just any of them can be a cause of something; a variable can be a cause only if it possible to "intervene" on it—by which they mean, in part, to cause it to take another value. One might doubt that it is possible to cause the raccoon in my neighborhood to become a member of a different species.

background condition. And in fact *O*'s having value 1 satisfies the condition stated earlier for being a cause of *L*'s having value 1—even though it is not a cause.

A causal model is often accompanied by a directed graph, where the nodes are the model's variables and an arrow is drawn from (the node for) *V* to *W* iff *V* appears on the right-hand side of the equation for *W*. I'll draw these graphs with dotted lines so they won't be confused with neuron diagrams. The graph for our model is in figure 7.

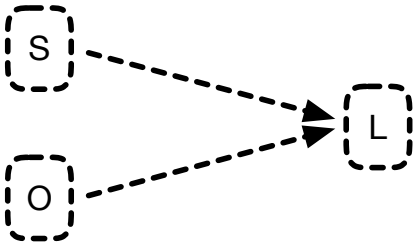

Figure 7

The roles of *O* and *S* in this graph, like their roles in the model itself, are symmetric. But, as I have argued, the presence of oxygen, and so *O*, is only indirectly relevant to the lighting, and so to *L*. To depict indirect relevance, we should draw an arrow, not from *O* to *L*, but from *O* to the arrow from *S* to *L*, as in figure 8, because what *O* does is modulate whether and in what way *L* depends on *S*.

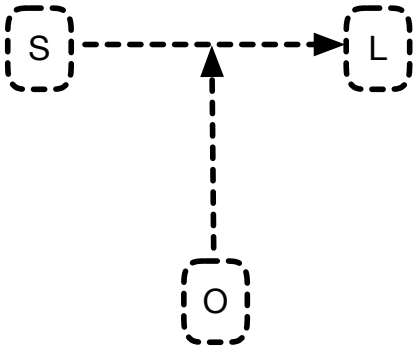

Figure 8

If we could depict causal models this way, direct and indirect modes of causal relevance would look different, and that would be an improvement.[27] But limitations would remain—there

[27]I mean: an improvement if you want a metaphysically clear representation of a situation's

would still be no way to distinguish variables that represent background conditions from those which represent enablers. For all that figure 8 shows, O could be either.

What about the causal models themselves, the sets of variables and equations? Incorporating into them a distinction between direct and indirect relevance is not so easy. One could complicate the models by designating some variables as modulator variables, and then determining the value of an endogenous variable using a function from modulator variables to a function from (directly relevant) variables to values. For example, the match-oxygen system would need a function from the modulator variable O to equations showing how S depends on L:

$$O = 1 \rightarrow L := S,$$

$$O = 0 \rightarrow L := 0.$$

This system is mathematically equivalent to the earlier equation, $L := SO$. But we are doing metaphysics, and the systems are not metaphysically equivalent. Only the second makes explicit that O is indirectly relevant to L.

10 Direct vs indirect relevance in scientific contexts

James Woodward, who favors a structural equations approach to causation, takes as his entry into thinking about causation the idea that causes are things "potentially exploitable for purposes of manipulation and control" (2003, v). Manipulating whether the match is struck and manipulating whether oxygen is in the room are both ways to control whether the match lights. If you think that the key to causation is manipulability then the striking and the presence of oxygen look to be on a par. If you think that the key to causation is manipulability then you may doubt that there is a metaphysically interesting distinction between causes and background conditions (and other indirect modes of relevance). You may welcome a formalism that makes that distinction invisible.

causal structure. A referee pointed out that, in another way, this change is not an improvement: it would make causal models useless for the causal discovery algorithms scientists have produced. But it's the first perspective that I care about here.

I think you should not welcome it, because the distinction between being a cause and being indirectly relevant to the causing is mirrored by a distinction between direct and indirect ways to manipulate or control something. Manipulating whether a match is struck is a direct way to control whether the match lights. Manipulating whether oxygen is present is an indirect way: what it controls is the way in which whether the match lights depends on whether it is struck.

Woodward might reply that this direct/indirect distinction, even if real, is one only metaphysicians care about, and so insofar as one wants a theory of causation that incorporates only distinctions that are important in science, one will ignore it. But the direct/indirect relevance distinction is important in science, as a few examples can establish.

In *The Selfish Gene* Richard Dawkins briefly takes up the question of why we die of old age, and (citing Peter Medawar) observes that natural selection will weed out genes that tend to harm or kill us when we are young, while leaving in place genes that tend to harm or kill us after we have reproduced. He observers further that (if this is right)

> we could try to 'fool' genes into thinking that the body they are sitting in is younger than it really is. In practice this would mean identifying changes in the internal chemical environment of a body that take place during aging. Any of these could be the 'cues' that 'turn on' late-acting lethal genes. By simulating the superficial chemical properties of a young body it might be possible to prevent the turning on of late-acting deleterious genes. The interesting point is that chemical signals of old age need not in any normal sense be deleterious in themselves. For instance, suppose that it incidentally happens to be a fact that a substance S is more concentrated in the bodies of old individuals than of young individuals. S in itself might be quite harmless, perhaps some substance in food which accumulates in the body over time. But automatically, any gene that just happened to exert a deleterious effect in the presence of S, but which otherwise had a good effect, would be positively selected in the gene pool, and would in effect be a gene 'for' dying of old age. The cure would simply be to remove S from the body.

> What is revolutionary about this idea is that S itself is only a 'label' for old age. Any doctor who noticed that high concentrations of S tended to lead to death, would probably think of S as a kind of poison, and would rack his brains to find a direct causal link between S and bodily malfunctioning. But in the case of our hypothetical example, he might be wasting his time! (1976, 41-42)

Dawkins thinks it is important and interesting, both to pure science and to medicine, that the presence of S might be, not a cause of bodily malfunction, but instead a context in which something else, lethal genes, have a "deleterious" effect.

Switching from popular to more esoteric science, endocrinologists say that some hormones in some contexts have what they call a "permissive effect." It turns out that to have a permissive effect is really to be an enabler. One textbook gives this definition:

> hormones may regulate receptors for other hormones. Estrogens increase the number of uterine receptors for progestins. Such an effect, in which one hormone induces production of receptors for a second hormone, or otherwise brings about the conditions necessary for the second hormone to be effective, is called a permissive effect. (Nelson 1995, 76)

Brings about the conditions necessary for the second hormone to be effective. That's almost exactly my definition of enabling. Deciding whether a hormone has a permissive effect is important in endocrinology, and permissive effects are common. Thyroid hormone is "essential for growth but is not itself directly responsible for promoting growth. It plays a permissive role"—note that this passage labels permissive effects as indirect (Sherwood 2015, 658). And in milk production the main cause is prolactin, which "acts on the alveolar epithelium to promote secretion of milk to replace the ejected milk," but "at least four other hormones are essential for their permissive role in ongoing milk production: cortisol, insulin, parathyroid hormone, and growth hormone." (Sherwood 2015, 769).

So scientists care about direct vs indirect relevance in their thinking about particular cases. Scientists also care in the abstract. Discussions of statistical methods emphasize the importance

of the notion of a "moderator variable," which is defined as a variable that "affects the direction and/or strength of the relation between an independent or predictor variable and a dependent or criterion variable" (Baron and Kenny 1986, 1174). Moderator variables are indirectly relevant. One textbook's "conceptual diagram" for moderation is, in fact, exactly my figure 8, showing one variable influencing the relation between two others (Hayes 2013, 209). How important is the notion of a moderator variable? That textbook's author writes that "moderation analysis" is one of "the more widely used statistical methods in the social, behavioral, and health sciences" (Hayes 2013, viii).[28]

11 Counterfactual dependence

I have distinguished between background conditions and enablers—an enabler causes the onset of a background condition. When I turn a knob to release oxygen from a tank, and fill the room where you're about to strike a match with oxygen, my turning of the knob is an enabler relative to the striking and the lighting. Enablers are events. Now if I hadn't turned the knob to release the oxygen, the match wouldn't have lit. If neuron C hadn't fired, in my neuron diagram of enabling, neuron B wouldn't have fired. These counterfactuals are still true. And isn't counterfactual dependence (between distinct events) sufficient for causation? It's sufficiency is after all enshrined in David Lewis's original counterfactual theory of causation (Lewis 1986a) and many of its descendants; the major problem its descendants face is the existence of cases where counterfactual dependence is not necessary.

The answer is no, counterfactual dependence is not sufficient for causation. The "Dependence Thesis" is false. Ned Hall has done more than anyone to blunt its appeal. Dependence, he argues, is inconsistent with the claim that causation is transitive, with the claim that causation is local in space and time, and with the claim that causation is intrinsic (roughly: a duplicate of a

[28]To be fair to Hayes, he also thinks (in my terms) that the direct/indirect relevance distinction is in how we think, not out in the world; conceptualizing X as a cause and M as a moderator, and conceptualizing M as a cause and X as a moderator, are both correct (Hayes 2013, 217). Needless to say, I disagree.

causal process connecting C to E is also a causal process). So you can't accept any of those and also Dependence; Hall argues that it is better to give up Dependence than the other three (Hall 2000; Hall 2004). Hall's arguments aside (well not entirely, since locality has come up in my argument too), I think that the appeal of Dependence fades as you explore the distinction between enabling and causing.[29] If C is a cause of E and N enables C to cause E, it may be (if we set aside redundant causation) that E counterfactually depends on both C and N. But the reasons why the counterfactuals are true are so different in the two cases! With the case of N (the enabler) there are all the complications about the state whose onset N causes, and that state's being a reason why C is cause; all these factor in to why the counterfactual is true. Not so in the case of the counterfactual about C. These differences more than justify jettisoning the idea that counterfactual dependence is sufficient for causation.

12 Two kinds of causation?; and, who cares?

I said that Hall argues against Dependence, but that's not quite right. His considered view is (or was) a version of causal pluralism: there are two concepts of causation, and Dependence is true only of the less central one. The central concept is the one that is transitive, local, and intrinsic. My view is that what Hall calls the central concept is the only concept, and that the one he says is characterized by Dependence is not a concept of causation at all. But I think that the list I have been compiling, of the varieties of causal relevance, can help someone with Hall's view. What is this second concept of causation, the one characterized by Dependence? What makes it tick? Can we get inside the way it works? What exactly is its relation to the central concept? They should answer that (what they call) the less central concept is a gerrymandered disjunction: to fall under that concept is to be either a cause or a background condition or an enabler.

But enough aid and comfort for the enemy. I've argued that indirect varieties of relevance are not causation. I reject Hall's causal pluralism; I deny that there is any kind of causation of which Dependence is true. But it may be that some version of causal pluralism can be defended without

[29]Lombard (1990) also thinks enabling falsifies Dependence.

confusing causation with other varieties of relevance. An example might[30] be Hitchcock's version: he holds that one way to be a cause is to be a "component cause" (roughly, cause along a particular causal path) and another is to be a "net cause" (cause when all paths are considered) (Hitchcock 2001b, 2003; a similar distinction appears in Woodward 2003; Pearl 2000 is the common cause). I insist that causal relevance comes in many kinds, not that causation itself does not.

Some people are okay with the idea that there is more than one variety of causal relevance, but not with my claim that only one of them is true causation—not because they disagree, but because they can't see why it matters. What real mistake is someone making who agrees that there are important differences between the presence of oxygen and the striking of the match, but insists that both are causes? I find this question frustrating, though my answer will surely frustrate those who ask it. The mistake is that their answer to *which things are causes?* is wrong, and that a wrong answer to this will inevitably lead to a wrong answer to *what is causation?*. And *what is causation?* is an important question. It's on the list of questions that are mandatory for a complete metaphysics. So it's bad to get the answer wrong.

I know people who won't accept that *what is causation?* is important until I tell them how answering it will help answer other important questions. While I think their demand can be satisfied (the examples are as usual: causation is connected to explanation and responsibility among other notions), I also think they demand too much, in this case and in general. *No question is important unless answering it helps answer some (distinct) important question* is obviously a vicious regress, entailing that no question is important. Some questions must be important for their own sake. *What*

[30]Why "might"? Hitchcock motivates his distinction with cases where something is a "positive" cause in one way and a "negative" cause in another, as in the philosophically-famous example of birth control pills, which are a negative net cause of thrombosis, but a positive component cause, along a path that excludes its effects on pregnancy. But (I say) there is no negative causation, only prevention (see section 15), which is a different variety of relevance. Also, my focus is the causation of particular events by particular events (so-called "token causation"), and so the question here is whether pluralism about token causation is correct; but the birth control pill example most directly supports pluralism about "type" causation ("birth control pills are a component cause of thrombosis" is a type-causal claim). It's doubtful this example supports token causation pluralism, since, as Hitchcock admits (2001b, 375), no particular pill-taking event can both cause someone to get a blood clot, and also prevent them from getting a blood clot.

is causation? is one of them.

13 Basic vs derived varieties

So far I have identified three varieties of causal relevance:

Name	Relation to Causation
Cause	—
Background Condition	Explains why C caused E rather than something else
Ennobler	Explains why C rather than something else caused E
Enabler	Cause of the onset of a background condition

In this list there is a pleasing symmetry between background conditions and ennoblers. Each explains why C caused E, but the explanations target different contrasts. Background conditions target why C had one effect rather than another—wiggle the background conditions, and you change what effects C has. Ennoblers target why E had one cause rather than another—wiggle the ennoblers, and you change which of the events preceding E earns the status of cause. The category of an enabler, however, is a dangler. It doesn't really belong with the other three. *Cause, background condition,* and *ennobler* are what we might call basic mode of causal relevance. But *enabler* is a derived mode, got by applying *cause* to *background condition*.

Now if applying *cause* to *background condition* generates a derived mode, then presumably combining any two (basic or derived) modes generates a derived mode. Something might be a cause of an ennobler, or a background condition to something's causing an enabler. Once we contemplate this ever-expanding hierarchy of derived modes, we might wonder whether the expansion always goes on forever, or whether instead some combinations of modes "collapse" into lower modes. If $M1, ..., Mn$ are modes of causal relevance, and so are $R1, \ldots, Rm$, does it ever happen that $Mn...M1$ (the result of applying $M2$ to $M1$, and then $M3$ to the result, and so on) is identical to $Rm \ldots R1$? Enablers of E are not (thereby) causes of E, but are causes of facts that ennoble C into a

cause of E thereby causes of E? While I doubt that collapse happens very often, these questions are hard, and I will mostly avoid them, except for one that has drawn a lot of attention: the question of whether preventers of preventers are causes. But I can't discuss double prevention until I discuss single prevention. Prevention is a kind of negative causation, and so a kind of causal relevance. But where does it fit on the list?

14 Transitivity

Having raised questions about prevention, I want to postpone them. First I want to use the distinctions made so far to resolve part of the debate over whether causation is transitive: whether, necessarily, if C is a cause of E and E is a cause of F, C is a cause of F.

Transitivity is threatened by switching examples. A train track starts at station A and divides into two parallel tracks which reconverge just before station B (figure 9). Back at station A a train departs. Before it reaches the division point I pull the lever that sets the junction to send the train down the lower track. It does so, and later arrives at station B. The argument goes: my pulling the lever causes the train's being on, or moving along, the lower track; its moving along the lower track causes its arrival at B; but my pulling the lever does not cause its arrival, since it would have arrived at B anyway (indeed, the case can be set up so that it would have arrived at the same time). If all that is right, then this is a counterexample to transitivity.

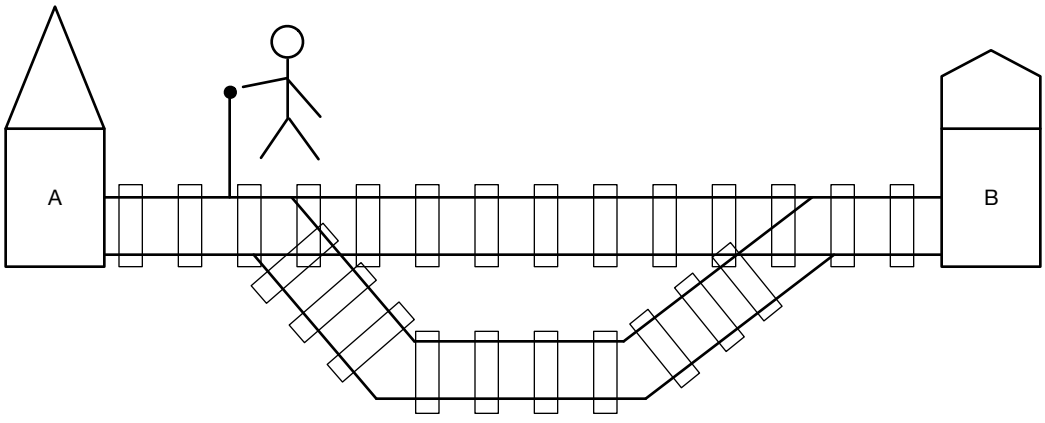

Figure 9

I shall argue that switches are not counterexamples to transitivity. Now Hall (2000) claims that all problem cases for transitivity either involve short circuits, or switching, or both. He also argues that short circuits are only counterexamples to transitivity if there is causation by double-prevention.[31] Since I deny that the double-prevention in short-circuits is causation (section 20), showing that switches are not counterexamples makes for a full defense of transitivity. Hall argues for this too, but, as I will explain, my argument that switches are not counterexamples differ from his.

If causation is transitive, then either (i) my pulling of the lever is, after all, a cause of the train's arrival; or (ii) my pulling of the lever is not, after all, a cause of the train's moving on the lower track; or (iii) the train's moving on the lower track is not, after all, a cause of the train's arrival at B. Claim (iii) is clearly false.[32] Hall argues at length that option (i) is the way to go (2000). But that's the wrong solution. Really it is (ii) that is correct, and the concept of an enabler puts us in a position to see this. My pulling of the lever is an enabler: it enables the train's departure to cause the train's motion along the lower track. It puts in place a state—one consisting in the tracks at the junction pointing toward the lower track—that is a background condition to the departure's causing the motion on the lower track. But something that enables the departure to cause the motion is not itself a cause of the motion. If you lack the concept of an enabler it is easy to think that, since the pulling of the lever is obviously causally relevant to the train's motion on the lower track, it must be a cause of the train's motion. But the existence of other varieties of causal relevance shows this to be bad reasoning.

Hall's arguments that my pulling of the lever is, after all, a cause of the train's arrival are faulty for similar reasons. Here is one of them:

[31] A canonical short circuit: "Billy sees Suzy about to throw a water balloon at her neighbor's dog. He runs to try to stop her, but trips over a tree root and so fails. Suzy, totally oblivious to him, throws the water balloon at the dog. It yelps" (Hall 2000, 201). Billy's beginning to run is a cause of his tripping, and if double-preventers are causes, his tripping is a cause of yelping, but the running is not a cause of the yelping.

[32] On a view like Paul's (2000), (iii) is true: it is the train's moving toward B, an event that could have happened on either track, that is a cause of the arrival. Hall (2000) argues that this response does not generalize to all switching cases.

> Notice, first, that among the (highly nonsalient) causes of the train's arrival is the presence of a certain section S of the track down which the train travels, during the time that the train is, in fact, on S. And among the causes of the presence of S at this time is the presence of S a day earlier. So we should say: among the causes of the arrival is the presence of S a day earlier. That is certainly odd, but it is odd in the familiar way to which writers on causation have grown accustomed: whatever the exact nature of the causal concept we are trying to analyze, it surely picks out a very permissive relation, one which does not distinguish between events that we would normally single out as causes and events we would normally ignore because their causal relationship to the effect in question is too boring or obvious to be worth mentioning, or is easily hidden from view as part of the "background conditions." The presence of S a day earlier is like that.
>
> But the relationship that the switching event bears to the arrival is just like the relationship that the earlier presence of S bears to the arrival: for the relationship that the switching event bears to the setting of the switch as the train passes over it is just like the relationship that the earlier presence of S bears to the later presence of S, and the relationship the setting of the switch as the train passes over it bears to the arrival is just like the relationship the later presence of S bears to the arrival. [...] the switching event causes the setting of the switch, which interacts with the passing train in the same way as the mere presence of a piece of track interacts with the passing train, and so on. (2000, 207-208)

Yes, absolutely, this is almost exactly fifty-percent right. Hall draws an analogy between the role the switching (the pulling of the lever) plays in causing the train's arrival, and the role the presence of section S of the track plays (figure 10). It is a good analogy, but it works against the truth of (i), not in its favor. The switching and the presence of S are similar, not because both are causes, but because neither is. The presence of S is a state that has obtained ever since S was installed. It obtained when the train passed over S, and it obtained a day earlier. And that state

was a background condition to the train's departure's causing the train's arrival. It is part of why the departure caused the arrival. But that does not make the presence of S a cause of the arrival; a background condition to C's causing E is not a cause of E. So if we are to look at the role the presence of S plays when figuring out the role that the junction's pointing toward the lower track plays, we will want to say that it too is a background condition (to the train's departure's causing the train's motion on the lower track). The same goes if we are to look at the role that the presence of S a day earlier plays when figuring out the role that the pulling of the lever plays. Here the analogy is not perfect: the pulling of the lever causes the junction to begin to point toward the lower track, while the presence of S a day earlier does not cause the presence of S today. Still, neither is a cause. (The earlier presence of S is the same state as the presence of S today, while the pulling of the lever is an enabler.)

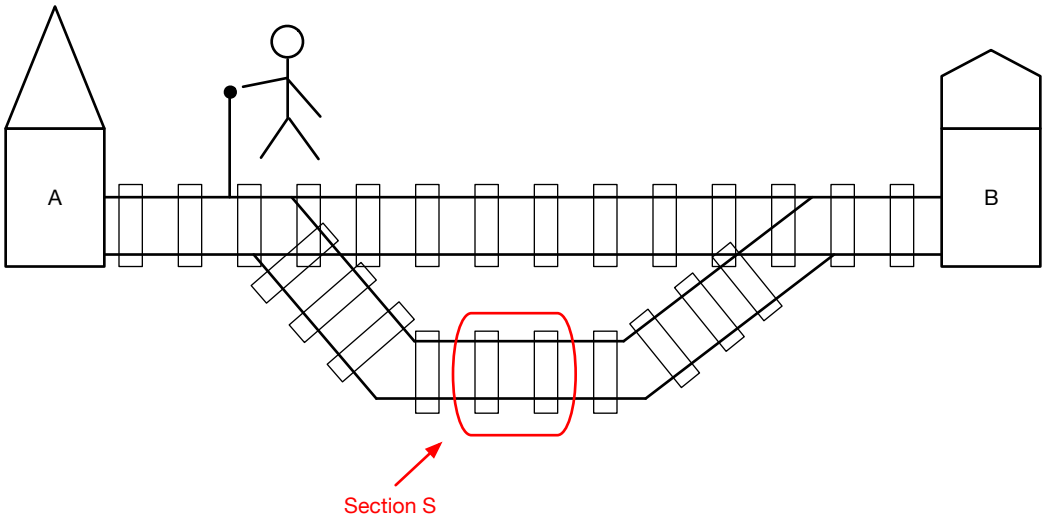

Figure 10

Hall is willing to call the presence of S a background condition, but he doesn't think that shows it is not a cause; he thinks that background conditions are non-salient causes. And he thinks they are causes because he accepts a very liberal theory of events, one that counts the presence of S as an event. If that were an event, then yes, it would be hard to see how it differed from the train's departure. But that liberal theory is false; the presence of S is not an event. Events, again, go with doings, but in being present S is not doing anything.

Hall also uses a "subtraction test" to argue that the pulling of the lever is a cause of the train's arrival (2000, 206). Imagine that a portion of the upper track, the track the train never passes over, had not been there. In that scenario the pulling of the lever would have been a cause of the train's arrival (right?). But whether the pulling is a cause shouldn't depend on what is going on with a bit of track that the train never gets near. So even when the upper track is intact, the pulling is a cause. Right: whether the pulling is a cause shouldn't depend on something far away from where the train ever was. But the right conclusion is not that the pulling is a cause even when the track is present, but that the pulling is not a cause even when the track is absent. In a scenario where part of the upper track is missing, the pulling of the lever to set the junction to send the train to the lower track is (still) not a cause; it is still (just) an enabler. Pulling the lever puts into place a state that is a background condition to the train's departure causing its arrival. Pulling the lever, or leaving it in place, is a way of manipulating what states obtain, and so manipulating the background conditions. The difference is, with the upper track in place, either way, pull the lever or do nothing, a state will be in place that will determine that the train's departure from A causes its arrival at B. Not the same state: pull and one state will obtain, do nothing and another will. But either state determines the departure to have the same (ultimate) effect, the arrival at B. When a piece of the upper track is missing, on the other hand, manipulating the background conditions matters to the departure's ultimate effect: pull the lever, and the departure causes the train to arrive at B; leave the lever where it is, and the departure causes the train to derail when it hits the broken track. But, again, neither enabling the departure to cause the arrival, nor enabling it to cause a derailment, is the same as causing the arrival, or a derailment.[33]

The way neuron diagrams for switches are sometimes drawn makes it hard to appreciate that they are enablers rather than causes, for reasons I've already mentioned: the invisibility of background conditions in those diagrams. Paul and Hall's canonical neuron diagram for switching is drawn as in figure 11 (2013, 232). The relation the firing of C bears to the firing of B does not appear to differ in any way from the relation the firing of A bears. But of course it does: the firing

[33]Similar remarks apply to Hall's "variation test" argument that pulling the lever is a cause of the arrival (2000, 207).

of C puts in place a state that determines which axon is excited when B fires. The firing of A does something different: it just causes B to fire. A more perspicuous diagram is in figure 12. We may imagine, fancifully, that which axon is excited when B fires depends on whether oxygen is present around B. This diagram makes clear the different roles that A and C play; it makes clear that C is an enabler.[34]

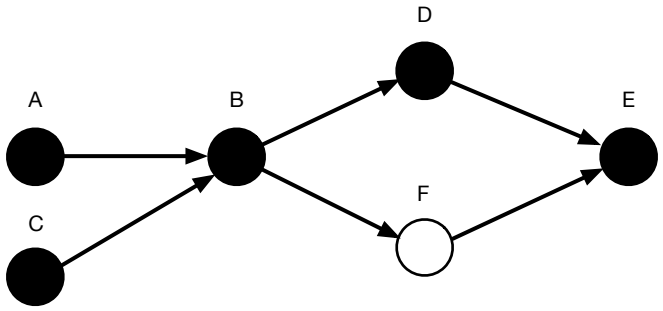

Figure 11

15 Prevention

A fourth basic way to be causally relevant is to prevent something. If I douse your match with water and then you strike it and it fizzles but does not light, I have prevented your match from lighting. Similarly if I suck the oxygen out of the room before you strike. Preventing is, in a way, a kind of negative of causing, though we will need to get clearer on the way in which this is so.

Preventing is different from the varieties we have already seen. If something exhibits any of the earlier kinds of casual relevance, then it is relevant to a causing. But prevention is not. If I

[34]A referee asked whether switches are essentially enablers. Suppose Jupiter shakes an asteroid loose and it heads toward the earth, but superman punches it away; but then, since he miscalculated, it bounces off the moon and hits the earth anyway. Superman's punch didn't enable Jupiter's pull to cause the asteroid to hit the moon. It was itself a cause of the asteroid's hitting the moon. But wasn't it also a switch with respect to the asteroid's path? Since the moon bounce is a cause of the asteroid's hitting the earth, transitivity is true only if superman's punch is a cause of the asteroid's hitting the earth—but isn't this the sort of claim I've been trying to deny? I admit to being troubled by this question and unsure what to say. My current view is that superman's punch is a cause of the earth impact, and also that the reasons for denying this are much weaker than the reasons for denying that switching the train track is a cause of the train's ultimate arrival.

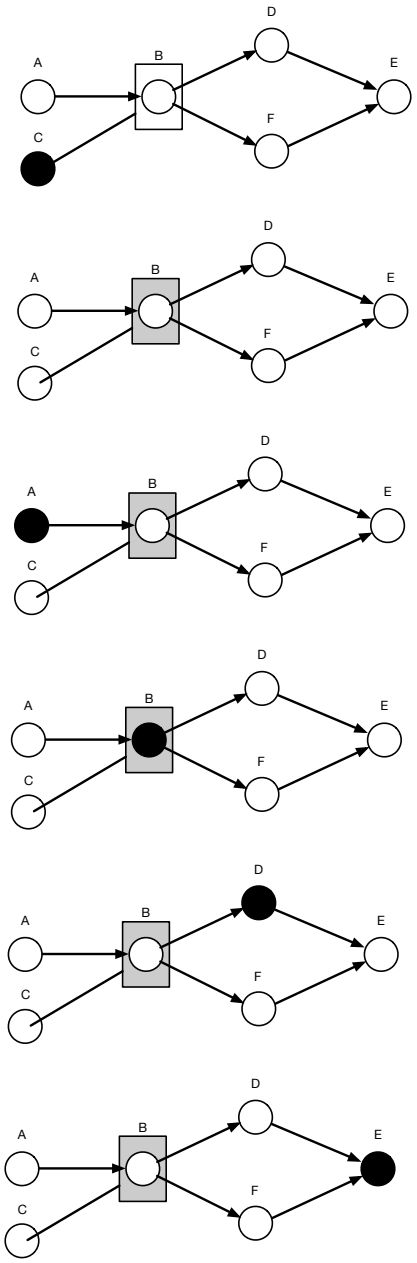

Figure 12

prevent the match from lighting, what I do is not relevant to the causing of the lighting, because, obviously, there was no lighting and no causing of a lighting. But if I prevent the match from lighting, what I do is relevant to something; it is relevant to your failure to cause the match to light. You failed because, although you did something that normally causes matches to light, namely strike the match; although you did something that started a process that was heading toward a

lighting of a match, as a result of my dousing it that process fizzled out before it could get there. One way in which preventing is a kind of negative of causing is that it is relevant to a failure to cause, rather than a success.

16 Against causing a negative: the analogy to undercutting defeat

One might object that preventing is after all a kind of causal relevance that we have already seen: it is a kind of causing, namely "causing a negative." Preventing the match from lighting is just causing the match not to light. This view says that preventing is bearing a "positive" relation (causation) to something negative ("not lighting"), while my view says that preventing is what it looks like on the surface: bearing a negative relation (prevention) to a positive thing (lighting).

One argument that preventing is not causing a negative goes by analogy. Preventing as a kind of causal relevance is analogous to undercutting defeat as a kind of evidential relevance. Here again is a canonical example of undercutting defeat: when facing a wall with my eyes open and inquiring into its color, if all I have to go on is how the wall looks, then that the wall looks red to me is evidence that it is red. But if I learn that the wall is illuminated by a red light, this new fact defeats my original evidence's support for the red hypothesis. It, as it were, prevents my first evidence from supporting the red hypothesis. Undercutting defeaters are certainly evidentially relevant. But undercutting defeaters are not evidence for a negative. That the wall is illuminated by a red light is not evidence that the wall is not red, or that it is some color other than red. So, given the analogy, preventing something from happening is not causing a negative.

17 Against causing a negative: there are no negative acts

I also have a direct argument that preventing is not causing a negative. Preventing is not causing a negative because it is not possible to cause a negative. It is not possible, for example, to cause a match not to light.[35]

[35]This is a metaphysical argument that preventing is not causing a negative. There is also psychological evidence that people do not equate preventing with causing a negative; see (Walsh and

I have been switching back and forth, since the beginning, between two ways to use "cause." On one use, the dominant one so far, it is events that cause and get caused. "The striking caused the lighting"—that's the first use. On the second use, it is things (that are not events) that cause and get caused. "*I* caused *the match* to light by striking it"—that's the second use. The cause is me, and I'm not an event, and the direct object of "cause" is "the match," and the match is not an event. If we label causation described the first way "event causation" then we can label causation described the second way "agent causation." Agent-causal claims are more complicated, because there are (or can be) two prepositional phrases following the direct object: I caused the match *(i) to light (ii) by striking it*. The two kinds of causal talk occupy different metaphysical levels, with (on my view) agent causation more basic, and event causation derived from it.[36]

Now notice that if I cause something to V, I cause it to do something; only when an activity verb goes in for "V" is the result something that can be true. Similarly, if I cause something to V by Zing, I cause it to V by doing something; only when an activity verb goes in for "Zing" is the result something that can be true. These observations provide the material for showing how event causation arises from agent causation: if X causes Y to V by Zing, then X and Y both did something, and so in virtue of their doing those things, two events happened, a Ving and a Zing; and the Ving is a cause (an event-cause) of the Zing. If I caused the match to light by striking it (that's agent causation), then my striking caused the lighting (event causation).

Let's operate at the level of agent causation. The argument that it is impossible to cause a negative is this. As I just claimed, causing is always causing something to do something. So causing a negative would be causing something to do a negative act. But there are no negative acts.

Why are there no negative acts? If there are any, there are lots of them. If there are any negative acts, then the operation that generates the negative of an act is defined everywhere; every act has a negative. But the act of waving my hand does not have a negative. If it did, "not waving my hand" would denote or express it, so it would make sense to say that one thing I am doing

Sloman 2011).

[36]For arguments that agent causation is more basic see (Lowe 2008), (Skow 2018), and (Baron-Schmitt MS).

right now is not waving my hand. But this does not make sense (and not on some philosophically sophisticated meaning of "not make sense"—it is literally ungrammatical).

Of course it does make sense to say that one thing I am not doing right now is waving my hand. But that is to mention a positive act (waving) that I am not performing, not to mention a negative act that I am performing.

Here are some more arguments against negative acts. Whenever something does something, it makes sense to ask how and where it did that thing. So if "not waving my hand" denotes or expresses a (negative) act, then whenever I am not waving my hand it makes sense to ask how and where. But these do not make sense. If I say I did not wave my hand just now you cannot meaningfully ask "How did you do it?" or "Did you do it quickly?"; you cannot meaningfully ask "Where did you do it?" or "Did you do it outside?" And while it is meaningful to say "I did not wave my hand quickly," this does not entail that I did something quickly. The meaningful interpretation here groups the words as "I did not [wave my hand quickly]"; if there were negative acts, there would be another meaningful interpretation corresponding to "I [[did not wave my hand] quickly]." But there is no such interpretation.[37] Finally, if X is something that can be done, then it always makes sense to ask how often I did it during some time period. So if there were negative acts, and so if "not wave my hand" denoted an act, then it would make sense to ask how often I did it. But neither "How often did you not wave your hand yesterday?" nor "I didn't wave my hand yesterday six times" (grouped as "I [[didn't wave my hand] six times yesterday]") make sense.[38]

Since there are no negative acts, it is impossible to cause something to perform one of them; it is impossible to cause a negative.

[37] A similar argument is in (Maienborn 2005, 310), but she is arguing that "I did not wave my hand quickly" does not entail that any event happened. For that argument, see below.

[38] See (Herweg 1991, 974); like (Maienborn 2005), he is arguing against negative events, not negative acts.

18 Against causing a negative: there are no negative events

When I say that you can't cause a negative I mean you can't cause something to perform a negative act. But most discussions of whether you can cause a negative ask whether you can cause a negative event. I have not yet said anything about whether there are negative events or, if there are, whether they can be caused. But my theory of the relation between events, acts, and doings implies answers to these questions. It implies that there are no negative events, and so no negative events around to get caused. For, again, events go with doings: if an event happens, that's because something did something. If that's what it takes for an event to happen, when does it take specifically for a negative event to happen? It has to be that the occurrence of a negative event requires that something perform a negative act. So if there is such a negative act as not-waving-one's-hand, and I perform it, then in virtue of this fact a negative event, a "not-waving of my hand," would happen. But, again, there aren't any negative acts to be done.

The ideas in the argument against negative acts can be used to construct a direct argument against negative events, one that does not depend on my theory of events.[39] Pronouns can refer to events that have not been explicitly introduced into the discourse, if a fact in virtue of which an event happened has been asserted. The sentence "A bomb exploded yesterday" does not itself mention any events, but since in virtue of its truth an event (an explosion) occurred, we can go on to say "...It was loud," where "it" refers to the explosion. So if there are negative events, then if a bomb (say, the biggest one in the local bomb shop) did not explode yesterday, in virtue of this fact a negative event occurred. So it should be okay to say "The bomb did not explode yesterday. It was quiet," where "it" refers to an event. But this is not okay. Similarly, "The bomb exploded yesterday. It happened in the remote desert" is fine, while "The bomb did not explode yesterday. It happened in the remote desert" can only evoke bewilderment. The conclusion, again: there are no negative events.

Now some philosophers, Woodward (2006) is a good example, resist denying that there are

[39] In fact, as I mentioned in note 37, I got the materials for the argument against negative acts from Maienborn's arguments against negative events (Maienborn 2005).

any negative events, on the ground that events can seem negative when conceptualized one way and positive when conceptualized another, when neither way of conceptualizing them is better than the other. Moreover, he says, some events that seem negative certainly exist and cause others. "[A]s an illustration of these points," he writes (note that Woodward uses "presence/absence" where I use "positive/negative event"—but he's clearly discussing the same distinction):

> consider death. Is it a presence or an absence? Looked at one way, it is about as clear a case of an absence (of life, of brain activity) as one could imagine. Nonetheless, we often treat death as a "positive" occurrence, and it is certainly not a mere "nothing" in the sense in which we might describe, say, an empty box as containing nothing. (Typically a body is present that is in a certain condition.) No one doubts that deaths can be effects (inquests investigate the causes of death), and it seems arbitrary to deny that deaths can be causes (of physiological changes to the body, of grief on the part of others, of the collapse of the empire) as well. (22)

Woodward doesn't come right out and say it, but his view seems to be that the question of whether death is a positive or a negative event is a defective question: "looked at one way" death is negative, looked at another it is positive, and neither way of looking is objectively correct, so all that can be said about whether death is negative is that it depends on how you look at it. And if this is right then a theory of causation should not care about which answer is correct. A theory of causation should not give positive events and negative events different statuses, or causal powers.

I don't think the question of whether death is a negative event is a bad question. I don't think it depends on how you look at it. And I know the (objectively correct) answer: death is a positive event. That's because something's death is an event that happens in virtue of the fact that it died. Dying is an act, it is something that something can do. And dying is a not a negative act. It is not a negative act because, as I argued above, there are no negative acts.

Now my argument against negative acts assumed that if there are negative acts then some of them can be denoted or expressed by phrases of the form "not Ving" (not lighting a match, not waving my hand). Someone might reject this premise, and say that while there are negative acts,

and that every act has a negative, none of them can be denoted or expressed by phrases of this form, and few can be denoted or expressed in English in any way. Dying is one of these rare few. This seems like a wild position to me. If negative acts are as abundant as positive ones, it would be important to be able to talk about them, and so a terrible defect in English that it does not provide a simple way to take an activity verb phrase V and generate a verb phrase denoting the negative of the act V denotes.

But the merits of this position in general don't matter. There are independent reasons to deny that dying is a negative act. If dying is a negative act, it is the negative of living. There doesn't seem to be any better candidate positive act for dying to be the negative of. But if dying is the negative of living, then if something is dying, it is doing so in virtue of not living. That's just what it means to call the act of dying the negative of the act of living. But it is false that if something is dying, it is doing so in virtue of not living. A person suffering through the end-stages of a deadly illness is dying, and is also living. They are doing both. And if they are doing both, it cannot be that they are dying in virtue of the fact that they are not living.[40]

I am not saying, note, that there is no relation between the process of dying and the process of living. Maybe something is dying if something is happening to it that will bring its life to an end. Defining death and dying in terms of life faces obstacles, but even if this definition is right, it would not make dying the negative of living.[41]

Of course whether one should recognize the existence of negative events is and should be a complicated affair. I've presented what one might call "language-y" arguments against them, but one might maintain that there are still quite important ways in which a metaphysical theory that postulates negative events is better than one that does not. Here is one way you might argue. If I douse your match with water just before you strike it, surely there's something causally relevant going on; it is very different from a situation in which the match just sits in a cool room, unmoved,

[40]Paul and Hall (2013, 179-80) reply to a similar argument by saying that both life and death are positive events, but do not have the connection between negative events and negative acts to appeal to in support.

[41]Feldman (1994) discusses the obstacles.

even if, in both cases, the match does not light. If we postulate negative events, the idea goes, we can capture the difference by saying that in the first case, I caused a negative event (the match's failure to light), while in the second case, nothing caused anything (relevant).

If there were no other good way to capture this difference other than by postulating negative events (and maybe also the negative acts needed to generate them), that would constitute at least some reason to believe in them. But there is another way to capture this difference. Instead of complicating our theory of acts and events, we can—as I am suggesting—complicate our catalogue of kinds of causal relevance. We can reject the thesis that preventing is causing a negative, and say that in the first but not the second case something prevented the match from lighting.[42]

19 Preventing is indirectly relevant

There are other problems with taking preventing to be causing a negative. Preventing is more like enabling than causing: it is a kind of indirect causal relevance. So taking preventing to be causing a negative makes it, wrongly, a kind of direct causal relevance. When I enable the match to light by filling the room with oxygen, I manipulate the conditions in the room so that they become conditions favorable to your striking having its usual effect. I make the conditions into background conditions to your striking's causing the lighting. That's a case of enabling; what is its opposite? A case of preventing.[43] When I douse your match with water, I manipulate the condition of the match so that it becomes a condition unfavorable to your striking having its usual effect. In one central kind of case at least, when X prevents Y from Zing, something has done something that would otherwise cause Y to Z; and what X has done is interfere with the process before it can achieve

[42]What I say here about postulating negative events applies also to saying that the basic kind of causation relates facts, so that "negative facts" like the fact that the match did not light could be effects (see Mellor 1998). Better to understand prevention as a species of causal relevance distinct from causation. (I think there are other problems with fact causation; see chapter 5 of Skow 2018.) A philosopher might agree that there are no negative events, but maintain that preventing is causing, just causing a positive event "negatively described"; see (Schaffer 2005) for one approach. For reasons of space I cannot discuss this view; see (Paul and Hall 2013, chapter 4) for criticism.

[43]Lombard (1990) also notices that preventing is an opposite of enabling: he calls preventing "disenabling" (202).

this, by manipulating the conditions that determine the course that process will follow.[44]

It becomes clearer that preventing is, like enabling, an indirect kind of causal relevance, when you look at standard neuron diagrams containing prevention. The simplest neuron diagram depicting prevention is in figure 13. There are three neurons—that's important. Neuron A is hooked up so that it can send neuron C a stimulatory signal. Neuron B is hooked up so that it can send C an inhibitory signal. The "neuron laws" say that if C is receiving an inhibitory signal, then it will not fire, even if it receives a stimulatory signal.[45] In the scenario this diagram depicts, A and B both fire, and C does not, and it does not fire because it is being inhibited by B's signal. That is: B prevents C from firing, by preventing A's signal from causing C to fire. Clearly, then, preventing is indirectly relevant: B's signal is relevant to whether A's signal can cause C to fire.

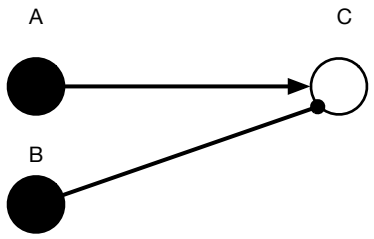

Figure 13

[44]Other philosophers who deny that preventing is causing a negative include Dowe (2000, 2001) and Moore (2009, 453): Dowe, because he thinks that a physical process must connect cause to effect, and such processes cannot include negative events; Moore, because he thinks that there are no negative events to be caused. Hitchcock (2007) is agnostic about whether preventing is causing, because the relation between a preventer and what it prevents is not, in his terms, "self-contained." While this phrase has a quite technical meaning in his theory, it is not far from the idea that preventing is an indirect kind of causal relevance. Dowe's counterfactual theory of prevention also in effect treats prevention as a kind of indirect relevance.

The psychological literature contains support for the idea that prevention is a kind of indirect relevance. Cheng (1997) presents evidence that subjects infer the existence of what she calls a "preventive" cause only when a (potential) "generative" cause is also present. Woodward (2006) cites her paper and asserts that "the operation of a preventing cause requires the presence of a generative cause but not vice-versa. Preventive causes prevent by interfering with or blocking generative causes" (29). (These authors assert not just that prevention is indirect but also that it is a kind of causation, while I am arguing that prevention's indirectness precludes it from being causation.) The psychological experiments on prevention in (Walsh and Sloman 2011) also suggest that subjects think of prevention as indirect.

[45]This is the standard way of understanding inhibitory signals, enshrined in the canon by David Lewis's presentation of the neuron laws (1986a, 200-1).

I spoke of one kind of preventing being, like enabling, a matter of manipulating the conditions: in the one case, to make them unfavorable to the effect, in the other, to make them favorable. That's what's going on here. The relevant condition is that of neuron C. If B does not fire, C is "receptive": it will fire if stimulated. But if B does fire, it makes C unreceptive: it will not fire if stimulated.

Since what is being manipulated here are conditions, that is states, figure 13 is misleading, since it is not drawn in a way that makes visible what states obtain when. We should re-draw it using our conventions for representing states. First we need to specify what the neuron looks like when receptive, when unreceptive. So let its wall be of normal thickness when it is receptive, and be extra thick when unreceptive. Then the sequence of events in figure 13 unfolds as in figure 14. This figure makes it even more clear that preventing is like enabling; it makes it even more clear that preventing is changing what states obtain that determine the effects of a given cause, in this case the firing of A. (I have B fire right before A, to make clear that the inhibitory signal gets to C before the stimulatory one.)

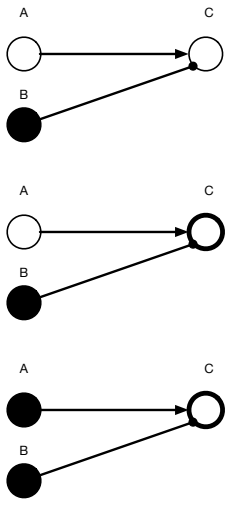

Figure 14

If preventing were a kind of direct causal relevance, as causation itself is, then we ought to be able to draw a two-neuron diagram containing prevention. But we cannot. Of course nothing stops us from writing down the diagram in figure 15. It has just neuron B and its inhibiting connection

to neuron C. B fires, and sends an inhibitory signal to C. But that's all: it is not true in this scenario that B prevents C from firing. A neuron that sends an inhibitory signal does not prevent anything, does not prevent the signaled neuron from firing, unless that neuron is also being sent a stimulatory signal.

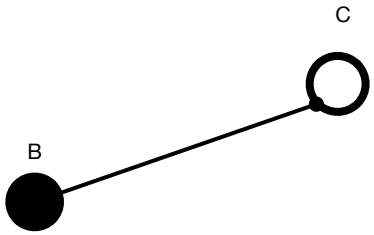

Figure 15

Of course you could imagine neurons that behave differently from these. What if neuron C, in figure 15, fires spontaneously, every five seconds, but B sends an inhibitory to C just before the five seconds are up, so that it doesn't fire? Then, yes, B prevented C from firing, even though there are only two neurons in the diagram. But this does not undermine my point. The prevention in this scenario is still indirect: B's signal prevents whatever mechanism in the neuron it is that recognizes that five seconds have passed from causing the neuron to fire. There is still a potential cause of the firing that is blocked, it is just not another neuron (and so is not represented in the diagram).

My thesis is really that in some cases preventing is the opposite of enabling. There is another variety of prevention. Enabling as I have defined it always involves manipulating the conditions that determine the course of a causal process: enabling steers the process toward an effect. In the examples so far preventing also involves manipulating the conditions, to steer the process away. But there are other ways to prevent one thing from causing another, besides manipulating the conditions. If you throw a baseball at a window, and I blast the baseball with my laser gun before it gets to the window, I prevent you (and the baseball) from breaking the window, but not by manipulating which states obtain. The same goes if I throw another baseball at yours, knocking it aside. That's not changing the states that determine the baseball's path; it is, in a more direct way, causing the baseball to take a different path.

20 Double prevention

Enabling, I said in section 13, is not a basic mode of causal relevance; instead it is causing the onset of a background condition. This observation raised the question of when one combination of varieties of relevance reduces to another; and I set the question aside. Now I want to bring it back. Does the combination *preventing a prevention* reduce to causation? Is double prevention causation?

Examples of double prevention became prominent in debates about causation largely through the efforts of Hall (2000; 2002; 2004) and Schaffer (2000). One of Hall's canonical examples is *Newt's Failed Mission*:

> It's a typical day during World War III, and the Good Guys have sent Bomber Billy on a mission to destroy an enemy target. Earlier in the day, one of the Bad Guys' fighter pilots—let it be Nasty Newt—takes off in his fighter plane on a routine patrolling mission. The Bad Guy ground-based radar defense system has just spotted Billy, and orders are about to be sent to Newt to go shoot him down, when a malfunction occurs in Newt's plane, and, sensitive thing that it is, it explodes. Billy's mission goes through completely undisturbed, and he destroys his target. But if the malfunction *hadn't* occurred, Newt would have shot down Billy, and the target would *not* have been destroyed. So the malfunction prevents something (Newt shooting down Billy) which would have prevented something else (Billy destroying the target) [...]. (Hall 2002, 277–78)

Figure 16 shows the events of the story. Hall's conclusion is complicated. There are two kinds of causation, he says, "production" and "dependence"; and the malfunction is not a productive-cause of the destruction of the target. Double prevention is not (productive) causation. I think there is just one kind of causation, though many varieties of causal relevance. But the relevant changes being made, I agree with Hall: the malfunction is not a cause of the destruction.

For more evidence that Newt-style double prevention is not causation, consider psycholog-

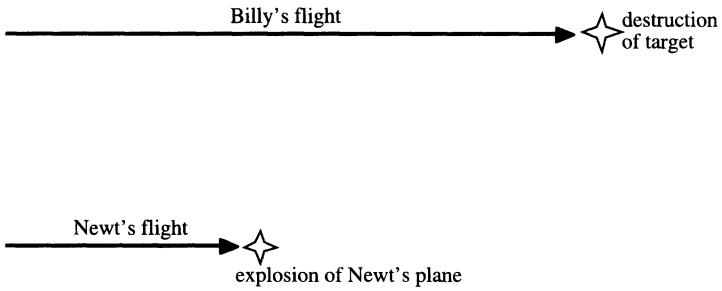

Figure 16

ical theories of personal identity. They say, with a lot of qualifications and clarifications, that a person P at a time t is identical to a person O at some earlier time iff P's psychological states are caused by O's psychological states. A common motive for these theories is the possibility of teletransportation. The person who appears on the planet after the transporter machine has read off Dr Spock's mental state and beamed that information down—that person is Dr Spock. But if double prevention is always causation, then these theories are not worth taking seriously. If I step into the transporter, which beams information about my mental state toward the planet, and that beam hits an asteroid, knocking it off coarse before it could block another beam of information, a beam that by coincidence was a duplicate of the beam the transporter sent but which was there by some freak statistical-mechanical fluctuation, and this second beam arrives on the planet and creates a duplicate of me, then, if double prevention is causation, the person on the planet's psychological states are caused by mine. My mental states prevent the asteroid from preventing the second beam from reaching the planet and creating my duplicate. But in this case the person on the planet is certainly not me.

Case closed? Definitely not. The clear verdict that these cases of double-prevention are not causation is matched by opposite but equally clear verdicts about other cases: some cases of double-prevention are quite obviously cases of causation. Cannonical here is Schaffer's description of how guns fire and the accompanying blueprint for a gun (figure 17):

> pulling the trigger causes the removal of the sear from the path of the spring, which causes the spring to uncoil, thereby compressing the gunpowder and causing an explo-

sion, which causes the bullet to fire. (2004, 199)

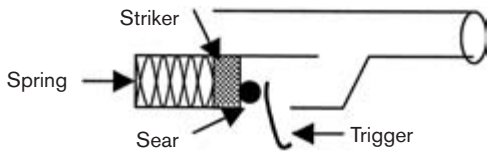

Figure 17

Pulling the trigger prevents the sear from preventing the spring from uncoiling. And pulling the trigger causes the firing of the gun.

Most treatments I am aware of treat all double prevention as of a piece: either it is always, or it is never, causation. But the cautious conclusion to draw from these examples is that double prevention sometimes is, and sometimes is not, causation.

This of course is unsatisfying if nothing can be said about when it is and when it isn't. And in fact there is an argument that nothing can be said. Newt's Mission and Schaffer's gun are, the idea goes, indistinguishable, suitably described. After all, both are represented by the same neuron diagram, and therefore there can be no causal difference between them. Well are the neuron diagrams really the same? The standard diagrams for the two are in figure 18 (adapted from Hall 2004, 242) and figure 19 (adapted from Schaffer 2000, 287). And these are the same, differing only in how they are labeled.[46]

But the examples are not indistinguishable. Double preventions can be divided into the causal and the non-causal. Here is how.

In Newt's Failed Mission, the malfunction on Newt's plane prevents Newt from intercepting Billy. Newt is a would-be preventer, stopped by the malfunction before he can interfere with Billy.

[46]The diagram in (Hall 2004) is not labeled with names for events in the story, but the intended labeling is clear. (The story is also slightly different from the one in Hall 2002, in irrelevant ways.) The diagram in (Schaffer 2000) has its neurons in different places, but my moving them does not change the meaning of the diagram. I've also added a neuron and two labels for clarity: "spring cocked" (in its second occurrence) and "spring uncoiling." If you are wondering about how to interpret the unlabeled white neuron in the diagram, stay tuned.

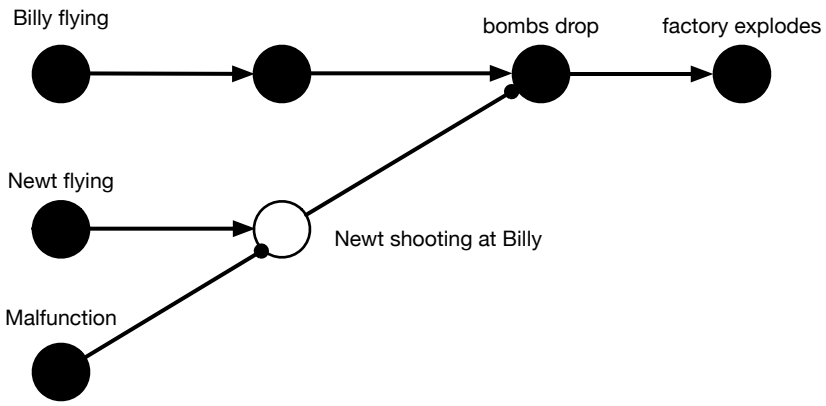

Figure 18

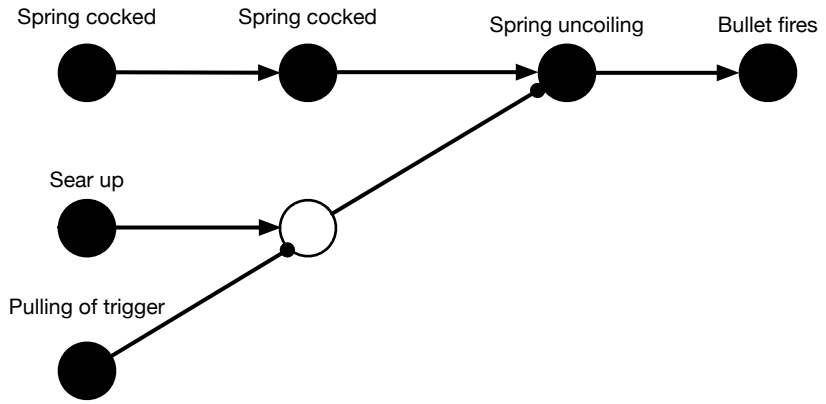

Figure 19

That makes it a case of what I will call blocking double prevention. In blocking double prevention, one process produces the final event, in this case the destruction of the target; and a second process that would have prevented it from doing that is blocked before it can even begin preventing it. At no time is Newt preventing Billy from doing anything. Newt's preventing is purely hypothetical: something he might have done, had the malfunction not happened.

Schaffer's gun, by contrast, exhibits what I will call interrupting double prevention. Pulling the trigger prevents the sear from preventing the spring from uncoiling, yes. But when the trigger does that, the sear has been preventing the spring from uncoiling for a long time. If the gun has been cocked for an hour before it is fired, then the sear has been preventing the spring from uncoiling for an hour. In interrupting double prevention, the process that produces the final event, in this case the spring uncoiling and smashing the striker into the back of the bullet casing, has been interfered

with, held back, by something else, for some stretch of time, before something, here the pulling of the trigger, interrupts this interference; it prevents that preventer from continuing to hold it back.

Interrupting double prevention is when one thing stops another from continuing to prevent something; blocking double prevention is when one thing prevents something from even beginning to prevent something. My thesis is that interrupting double prevention is causation and blocking double prevention is not.

I've extracted this thesis from two examples; let's look at more. First, Paul and Hall's two pillars (2013, 191). Two pillars are leaning against each other, holding each other up. You push one of the pillars away, and the other falls. This is double prevention: you prevent the first pillar from preventing the second pillar from falling. And it is causation: your shoving the first pillar is a cause of the falling of the second. My thesis agrees, since this is interrupting double prevention. The first pillar was already preventing the second from falling before you acted. When you shoved it, you stopped it from continuing to prevent this.

Next, Bennett's heavy rain example:

> There was heavy rain in April and electrical storms in the following two months; and in June the lightning took hold and started a forest fire. (1987, 373)

The rain is a double-preventer: it prevents the May lightning from preventing the June lightning from burning the forest. But the rain is not a cause of the fire—it takes a lot of philosophy to shake people's conviction in this. My thesis agrees. When the rain falls, the May lightning has not yet started preventing anything (it has not even happened yet!). The rain prevents the lightning from even beginning to prevent anything. This is a case of blocking double prevention.

Also a case of blocking is this example from Michael Moore:

> DP sees his old enemy V drowning in the ocean, which makes DP happy. Then DP spies a lifeguard, L, getting ready to save V. So DP ties up the lifeguard, and V drowns. DP has prevented [...] L from preventing V's death. But DP has not caused V's death. (2009, 62)

Right: because when DP acted, L had not yet begin to prevent V's death.

So my thesis sorts these cases of double-prevention the same way our intuitive judgments do. If you're worried that the diet of examples is too small, here are a few more. A bullet prevents someone's heart from preventing their brain from dying from oxygen starvation. Some calcium prevents some tropomyosin from preventing a muscle fiber from contracting. Both, Schaffer (2000) says, are cases of causation; and both are interrupting double prevention. In *E. coli,* lactose inactivates a protein that represses the genes that produce the enzymes needed to metabolize lactose. Woodward says that this is causation and that genetics texts agree (2006, 35); it is also interrupting double prevention.

Woodward (2006) offers an hypothesis (that competes with mine) about when we judge double prevention to be causation: double prevention seems causal to us when it is "stable," which means, roughly, that the counterfactual dependence between the final effect and the double preventer would still hold, under a wide range of changes in circumstance. Woodward's hypothesis entails that unstable interrupting double prevention will not look like causation. I submit that it does. For example: Suzy shoves pillar A, so that it stops supporting pillar B, which falls. The shove is a cause of the falling, even if many other shovers were ready to shove it instead if the room's temperature had been the least bit different. "Had Suzy not shoved, the pillar would not have fallen," while true, is not sensitive; even tiny changes to the background temperature would make it false.

Paul and Hall think that there is a strong case that double prevention is (always) causation, but look closely at their abstract description of double prevention:

> whenever two systems X and Y are stably interacting in such a way that the continued presence of X is inhibiting Y from behaving in a certain manner, then an event that "removes" X will thereby cause a change in Y, but do so by double prevention. (Paul and Hall 2013, 191)

This doesn't describe double prevention in general, only interrupting double prevention. Their generalization is right: that's always causation. Their mistake is thinking their description covers

all cases of double prevention.

What about arguments that double prevention is never causation that are driven by abstract principles about causation, rather than reflection on examples? When you look closely, those arguments actually target blocking double prevention, and not interrupting double prevention, and so support my thesis.

Hall offers two principled arguments (2000, 2002). First, causation is local in space and time, as long as we can ignore quantum mechanical weirdness. The spatial part of locality says that if C is a cause of E, then either C and E happen right next to each other, or there is a sequence of intermediaries C=X1,X2,...,Xn=E, each event a cause of the next, where Xi and Xi+1 happen right next to each other. But if double prevention is causation, it is non-local causation: the malfunction of Newt's plane does not happen right next to the explosion of the target, or any event involving Billy's bombs, or any event involving Billy's plane.[47]

Newt's mission is a case of blocking double prevention. And certainly the argument is right that, if blocking double prevention is causation, it is non-local causation. For in cases of blocking the prevention of the would-be preventer can be far away in space from the process leading to the final effect (Billy's flight, leading to the explosion of the factory). But blocking double prevention is not causation. What about interrupting double prevention: does it threaten locality? No. Now the would-be preventer cannot be far away from the processes leading to the final effect. In the case of Schaffer's gun, the preventing event, the trigger's pulling on the sear, happens where the sear is. And the sear is right next to the spring, and so right next to the uncoiling of the spring. And there is a clearly a sequence satisfying locality's demands from the uncoiling to the firing of the bullet. It is the same with the two pillars: the pillar holding up the other pillar is right next to it, in physical contact with it. The preventing event, someone's pushing of the first pillar, happens where that pillar is; which happens right next to the fall of the second pillar. Locality is satisfied.

Now Hall's second argument: causation is transitive, but if double prevention is causation,

[47] A natural thought is, if we accept the existence of omissions, then there is a local (omission-containing) sequence connecting the malfunction to the explosion. Hall argues that this thought is false.

there are counterexamples to transitivity. Suzy drops a tree branch above a bottle, and Billy, wanting to protect the bottle, starts running to intercept the branch, but trips, and falls, and fails; the bottle breaks. The tripping prevents Billy from preventing the branch from breaking the bottle. If double prevention is causation, the tripping is a cause of the breaking. But the running is a cause of the tripping and not also—as transitivity requires—a cause of the breaking. This is an example of a short circuit, and in general short circuits contain double prevention, and, if double prevention is causation, are counterexamples to transitivity.

But Billy's running and tripping is (another) case of blocking double prevention. The argument at most shows that, if blocking double prevention is causation, there are counterexamples to transitivity. But blocking double prevention is not causation. What about interrupting double prevention: does it threaten transitivity? No. Double preventers threaten transitivity when they appear in short circuits, but only blocking double preventers appear in short circuits. When a process X is on its way to causing an event E, a short circuit both launches a process P that threatens to prevent X from causing E, and also interferes with P so that it never succeeds. But only blocking double preventers never succeed.

Try turning Billy's attempt to stop Suzy's branch into a case of interrupting double prevention, while keeping it a counterexample to transitivity. It can't be done. Suppose Billy doesn't trip, but instead kicks a basketball, which caroms off a tree back toward him. He catches the branch for a moment, but then is hit by the ball and drops the branch, which continues on its way down and breaks the bottle. Billy's running is a cause of the ball's hitting him, which is a cause of the branch's falling (this is the interrupting double prevention), and the branch's falling is a cause of the bottle's breaking. For transitivity to fail, it would have to be that Billy's running is not a cause of the branch's falling (or the bottle's breaking). But it is. Interrupting double prevention does not threaten transitivity.

A third theoretical argument that double prevention is not causation is that if it is, then there can be causation by and of omissions; but there cannot (e.g., Dowe 2000, Moore 2009). Look again at the standard neuron diagram for double prevention (figure 20). Neuron C prevents Neuron

D from firing; if D had fired, it would have prevented E from firing. So if double prevention is causation, and so the firing of C is a cause of the firing of E, then it runs through the omission "neuron D's not firing."

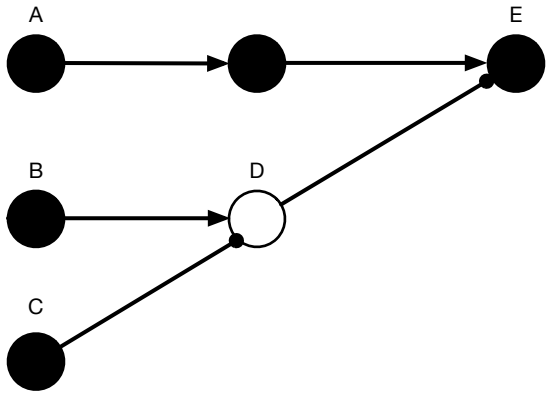

Figure 20

I accept the premise that there cannot be causation of or by omissions, that causation involves only "positive" events. But even with that premise the argument shows only that blocking double prevention is not causation. Only blocking double prevention runs through an omission. Interrupting double prevention does not. For example, the firing of a gun does not run through an omission. Before you pull the trigger, the sear is pushing against the spring, and a pushing is a positive event. When you pull the trigger, the sear stops pushing on the spring. This stopping pushing is then a cause of the spring's uncoiling. And a stopping pushing is a positive event; it is not an omission. I think this is intuitive—stopping pushing is a change from pushing to not pushing, and changes are positive events. It also follows from my theory of events, since "stop pushing" names something that can be done.

The argument that all double preventions are the same, causally speaking, was that all are aptly represented by the same neuron diagram. But this discussion of omissions shows that to be false. The diagram in figure 20 is only apt for blocking double prevention. In interrupting double prevention, the would-be preventer has already done some preventing, before the double-preventer stops it. But, as neuron diagrams are standardly read, neuron D in the diagram never sent any inhibitory signal to neuron E; it was never preventing E from firing.

Look back at how Schaffer labels his neuron diagram for the gun (figure 19, reprinted here).

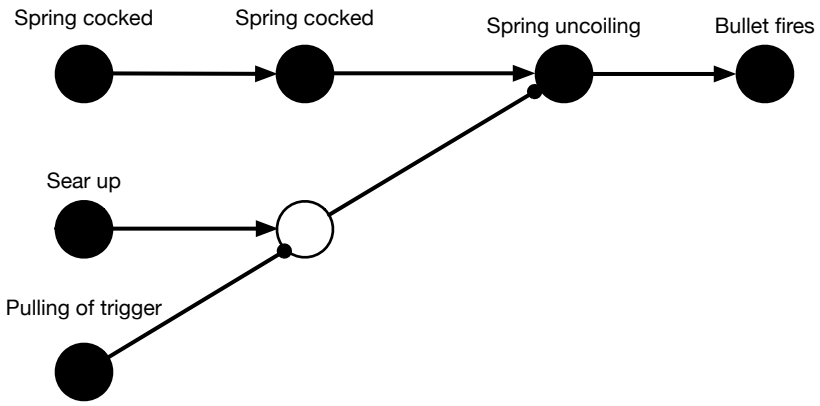

Figure 19

The neuron labeled "sear up" sends no inhibitory signal to any neuron representing the spring. That makes the diagram wrong: when the sear is up, it is preventing the spring for uncoiling.

So we can have our cake and eat it too. If A prevents B from preventing C from doing something, it might be that B had already started preventing C from doing something, when A stopped it; that is interrupting double prevention. Or it might be that B had not started; that is blocking double prevention. When an example of double prevention looks like causation, that's because it is; it is a case of interrupting double prevention, and they are cases of causation. And when an example of double prevention doesn't look like causation, that's because it isn't; it is a case of blocking double prevention, and they are not cases of causation. Theoretical arguments about double prevention and causation—that double prevention cannot be causation, because causation is transitive, and local, and relates only positive events—even if sound, are consistent with all of this, because those arguments are sound only if they concern blocking double prevention.

21 Enabling vs double prevention

One thread is left hanging. I have argued the enabling is not causing. Something that enables C to cause E is not (by virtue of that fact alone) a cause of E. And I said that to enable C to cause E is to put in place a state that is a background condition to C's causing E. These claims can appear to

conflict with the fact that double prevention is sometimes causation.

If a bomb explodes next to a dam, destroying it, and the water then rushes down the river and floods the village, then the explosion was a cause of the flood. This is causation by double prevention: the explosion prevents the dam from continuing to prevent the water from rushing down the river.[48] The dam was doing this by preventing gravity from accelerating the water downhill. But isn't this also a case of enabling? The explosion puts in place a state—an open space where the dam used to be—that is a background condition to gravity's pull causing the water to run down. If it is enabling, then enabling is causing, at least sometimes.

What all this shows, I think, is that enabling needs a more precise definition. True enabling, as a distinct variety of relevance, should be sliced off from causal double prevention. Enabling C to cause E is putting in place conditions in which C can cause E, where these conditions have not yet "become active" in the process C initiated. Removing an obstacle that is already holding back, or re-directing, this process, on the other hand, is not true enabling; it is causation by (interrupting) double prevention.

22 Conclusion

James Woodward wrote this about "what the notion [of a cause] is intended to contrast with":

> I believe that when 'cause' and cognate notions are used in the special sciences and in common sense contexts, the relevant contrast is very often with 'mere' correlations or associations. In particular, the underlying problematic is something like this: an investigator has observed some relationship of correlation or association among two or more variables, X and Y. What the investigator wants to know is whether this relationship is of such a character that it might be exploited for purposes of manipulation and control: if the investigator were to manipulate X (in the right way), would the correlation between X and Y continue to hold ...[or] is it instead the case that under

[48] Rickless (2011) and McGrath (2003) discuss cases like this.

> manipulation of X there would be no corresponding changes in Y [...]? In the former case, we think of X as causing Y, in the latter the connection between X and Y is non-causal, a mere correlation [...]. (2007, 71-72)

My thesis has been that this is wrong. If C is merely correlated with E, as for example the number of pool-drownings is correlated with the number of films Nicholas Cage stars in, then C is not in any way causally relevant to E.[49] But there are other alternatives to causation, and they are causally relevant: being a background condition, enabling, ennobling, and prevention. Recognizing these categories is important for its own sake, and has important benefits, including defusing threats to the transitivity of causation, and providing a clear view of when double prevention is and is not causation.[50]

References.

Bach, Emmon 1986. "The Algebra of Events." *Linguistics and Philosophy* 9: 5-16.

Baron, Reuben M. and David A. Kenny 1986. "The Moderator-Mediator Variable Distinction in Social Psychological Research: Conceptual, Strategic, and Statistical Considerations." *Journal of Personality and Social Psychology* 51: 1173–1182.

Baron-Schmitt, Nathaniel MS. "Thing Causation."

Bennett, Jonathan 1987. "Event Causation: the Counterfactual Analysis." *Philosophical Perspectives* 1: 367–386.

Cheng, Patricia W. 1997. "From Covariation to Causation: A Causal Power Theory." *Psychologi-*

[49] See `https://tylervigen.com/spurious-correlations`.

[50] Thanks to audiences at MIT, Harvard, NC State-Raleigh, and to Selim Berker, Ned Hall, Gideon Rosen, Jim Van Cleve, and an anonymous referee.

cal Review 104: 367–405.

Dancy, Jonathan 2004. *Ethics Without Principles.* Oxford University Press.

Davidson, Donald 1967. "The Logical Form of Action Sentences." In Nicholas Rescher, ed., *The Logic of Decision and Action.* University of Pittsburgh Press, 81–95.

Dawkins, Richard 1976. *The Selfish Gene.* Oxford University Press.

Dowdy, David 1979. *Word Meaning and Montague Grammar.* Kluwer.

Dowe, Phil 2000. *Physical Causation.* Cambridge University Press.

Dowe, Phil 2001. "A Counterfactual Theory of Prevention and 'Causation' by Omission." *Australasian Journal of Philosophy* 79: 216-226.

Feldman, Fred 1994. *Confrontations with the Reaper.* Oxford University Press.

Foot, Philippa 1978. "The Problem of Abortion and the Doctrine of Double Effect." In *Virtues and Vices*, University of California Edition?

Goldvarg, E. and P. N. Johnson-Laird 2001. "Naive causality: a mental model theory of causal meaning and reasoning." *Cognitive Science* 25: 565-610.

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

Hall, Ned 2002. "Non-Locality on the Cheap? A New Problem for Counterfactual Analyses of

Causation." *Nous* 36: 276–294.

Hall, Ned 2004. "Two Concepts of Causation." In *Causation and Counterfactuals.* John Collins, Ned Hall, and L. A. Paul (eds.) Oxford University Press, 225–276.

Hall, Ned 2007. "Structural equations and causation." *Philosophical Studies* 2007: 109-136.

Hayes, Andrew F. 2013. *Introduction to Mediation, Moderation, and Conditional Process Analysis.* The Guilford Press.

Herweg, Michael 1991. "Perfective and imperfective aspect and the theory of events and states." *Linguistics* 29: 969-1010.

Hitchcock, Christopher 2001a. "The Intransitivity of Causation Revealed in Equations and Graphs." *The Journal of Philosophy* 98: 273-299.

Hitchcock, Christopher 2001b. "A Tale of Two Effects." *Philosophical Review* 110: 361–396.

Hitchcock, Christopher 2003. "Of Humean Bondage." *British Journal for the Philosophy of Science* 54: 1–25.

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

Kenny, Antony 1963. *Action, Emotion and Will.* Routledge & Kegan Paul.

Kim, Jaegwon 1993. "Events as Property Exemplifications." In *Supervenience and Mind.* Cam-

bridge University Press, 33–52.

Landsburg, Steven 2012. *The Armchair Economist.* Revised edition. The Free Press.

Lewis, David 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61: 343–377.

Lewis, David 1986a. "Causation." In *Philosophical Papers Volume II.* Oxford University Press.

Lewis, David 1986b. "Events." In *Philosophical Papers Volume II.* Oxford University Press.

Lombard, Lawrence Brian 1990. "Causes, Enablers, and the Counterfactual Analysis." *Philosophical Studies* 59: 195-211.

Lowe, E. J. 2008. *Personal Agency: The Metaphysics of Mind and Action.* Oxford University Press.

Maienborn, Claudia 2001. "On the position and interpretation of locative modifiers." *Natural Language Semantics* 9: 191–240.

Maienborn, Claudia 2005. "On the limits of the Davidsonian approach: The case of copula sentences." *Theoretical Linguistics* 31: 275-316.

Maienborn, Claudia 2008. "On Davidsonian and Kimian States." In *Existence: Semantics and Syntax,* Ileana Comorovski and Klaus Von Heusinger, eds. Springer: 107-132.

McGrath, Sarah 2003. "Causation and the Making/Allowing Distinction." *Philosophical Studies*

114: 81–106.

Mellor, D. H. 1998. *The Facts of Causation.* Routledge.

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

Nelson, Randy J 1995. *An introduction to behavioral endocrinology.* Sinauer Associates.

Paul, L. A. 2000. "Aspect Causation." *Journal of Philosophy* 97: 235–5.

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

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

Pettit, Philip and Michael Smith 1990. "Backgrounding Desire." *Philosophical Review* 99: 565-592.

Plato 2002. *Five Dialogues.* G. M. A. Grube, translator, John M. Cooper, editor. Hackett.

Pollock, John and Joseph Cruz 1999. *Contemporary Theories of Knowledge.* Rowman & Littlefield.

Rickless, Samuel 2011. "The Moral Status of Enabling Harm." *Pacific Philosophical Quarterly* 92: 66–86.

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

Schaffer, Jonathan 2004. "Causes need not be Physically Connected to their Effects: The Case for Negative Causation." In *Contemporary Debates in the Philosophy of Science.* Christopher Hitchcock (ed.). Blackwell, 197–216.

Schaffer, Jonathan 2005. "Contrastive Causation." *The Philosophical Review* 114: 327–58.

Schroeder, Mark 2007. *Slaves of the Passions.* Oxford University Press.

Setiya, Kieran 2014. "What is a Reason to Act?" *Philosophical Studies* 167: 221–235.

Sherwood, Lauralee 2015. *Human physiology from cells to systems.* 9th edition. Cengage Learning.

Skow, Bradford 2018. *Causation, Explanation, and the Metaphysics of Aspect.* Oxford University Press.

Sloman, Steven, Aron K. Barbey and Jared M. Hotaling 2009. "A Causal Model Theory of the Meaning of *Cause, Enable,* and *Prevent. Cognitive Science* 33: 21-50.

Steward, Helen 1997. *The Ontology of Mind: Events, Processes, and States.* Oxford University Press.

Vendler, Zeno 1959. "Verbs and Times." *Philosophical Review* 66: 143–60.

Walsh, Claire R. and Steven A. Sloman 2011. "The Meaning of Cause and Prevent: The Role of Causal Mechanism." *Mind & Language* 26: 21–52.

Woodward, James 2003. *Making Things Happen.* Oxford University Press.

Woodward, James 2006. "Sensitive and Insensitive Causation." *Philosophical Review* 115: 1-50.

Woodward, James 2007. "Causation with a Human Face." In *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited.* Huw Price and Richard Corry (eds.). Oxford University Press, 66–107.

Yablo, Stephen 2010. "Advertisement for a Sketch of an Outline of a Prototheory of Causation." In *Things*.