



# MIT Open Access Articles

## *Are There Non-Causal Explanations (of Particular Events)?*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

<b>Citation</b>	Skow, Bradford. "Are There Non-Causal Explanations (of Particular Events)?" <i>British Journal for the Philosophy of Science</i> (2014) 65 (3): 445-467.
<b>As Published</b>	<a href="http://dx.doi.org/10.1093/bjps/axs047">http://dx.doi.org/10.1093/bjps/axs047</a>
<b>Publisher</b>	Oxford University Press
<b>Version</b>	Author's final manuscript
<b>Accessed</b>	Mon Mar 19 20:14:03 EDT 2018
<b>Citable Link</b>	<a href="http://hdl.handle.net/1721.1/73660">http://hdl.handle.net/1721.1/73660</a>
<b>Terms of Use</b>	Creative Commons Attribution-Noncommercial-Share Alike 3.0
<b>Detailed Terms</b>	<a href="http://creativecommons.org/licenses/by-nc-sa/3.0/">http://creativecommons.org/licenses/by-nc-sa/3.0/</a>

# Are There Non-Causal Explanations (of Particular Events)?\*

Bradford Skow

## Abstract

Philosophers have proposed many alleged examples of non-causal explanations of particular events. I discuss several well-known examples and argue that they fail to be non-causal.

1. Questions
2. Preliminaries
3. Explanations that Cite Causally Inert Entities
4. Explanations that Merely Cite Laws, I
5. Stellar Collapse
6. Explanations that Merely Cite Laws, II
7. A Final Example
8. Conclusion

## 1 Questions

What *is* explanation? What is its nature? One way to try to answer these questions is to gather together lots and lots of explanations and look for what they have in common.

Many explanations of particular events in ordinary life and in the sciences do have a common element. Many of them are causal explanations. This suggests an

---

\*Forthcoming in *The British Journal for the Philosophy of Science*.

hypothesis: maybe causal explanation is not just one kind of explanation. Maybe instead *what it is* to be an explanation is to be a causal explanation. Whatever set of features constitutes the nature of causal explanation also constitutes the nature of explanation (of events) more generally.

There is an easy way and a hard way to argue against this hypothesis. The hard way is to argue that, even if all explanations are causal, causation does not enter in to the correct account of what explanation *is*. Instead it (merely) follows from that account that (necessarily) all explanations of events are causal. The hard way is hard because it requires us to work with the (difficult) distinction between properties that something has merely of necessity and properties it has essentially, by its nature.

The easy way is the more common one: find an example of a non-causal explanation. Many philosophers have given examples of explanations and argued that those explanations are non-causal. While these examples may not have convinced everyone they do seem to have eroded philosophers' confidence in the thesis that all explanation (of events) is causal. In the past there were philosophers who explicitly affirmed this thesis: Salmon ([1984]) and Lewis ([1986]) are two examples. Philosophers affirm it more rarely today. James Woodward, for example, does not affirm it in *Making Things Happen* ([2003]) and Michael Strevens does not affirm it in *Depth* ([2009]) (these are two of the most influential recent books on explanation). Each leaves open the possibility of non-causal explanation.

I do not think that the standard examples of non-causal explanations of particular events are convincing. My plan for this paper is to discuss several examples and explain why they fail. The examples have been discussed before but there are new things to say about each of them. What I say here does not prove that there are no possible examples of non-causal explanations, but it does, I think, strengthen the case. Seeing why the examples I discuss fail will also help us better understand what causal explanation is and how it works.

## 2 Preliminaries

Although my topic is explanation my focus is only on explanations of particular events. My question is whether there are any explanations of particular events that are non-causal.<sup>1</sup> I will not be interested in explanations of laws of nature, for example, or explanations of mathematical truths.

But I am not interested in *all* explanations of particular events. I want to exclude one particular kind of explanation of events: ‘in-virtue-of’ explanations. Here is an example of one. Suppose someone asks ‘Why is the distance between A and B 5 meters?’ One (candidate) answer is: because the shortest path from A to B is 5 meters long. This is (or at least is trying to be) an explanation. It looks like an attempt to explain why some fact obtains by citing some other fact or facts that ‘ground’ the target fact, that are the ‘deeper’ facts ‘in virtue of which’ the target fact obtains.<sup>2</sup>

In-virtue-of explanations are obviously non-causal.<sup>3</sup> So the thesis that all ex-

---

<sup>1</sup>Some philosophers think that it is more philosophically fruitful to treat the causal relation as, at bottom, a relation between facts rather than events. (Bennett [1988] discusses the merits of the two approaches.) Anyone who accepts this view should accept a similar thesis about explanation: at bottom, the objects of explanation are facts, not events. If I were to take that view I would have to characterize my topic differently. But I do not think that anything in this paper depends on which of these views is correct, so I will continue to speak in terms of explanations of events.

<sup>2</sup>See (Rosen [2010]) for one discussion of in-virtue-of explanation.

Philosophers often aim to give in-virtue-of explanations, and think that a philosophical theory that gives better in-virtue-of explanations is thereby a better theory. But scientists also often aim to give in-virtue-of explanations. When a chemist explains why a pane of glass is fragile by describing its molecular structure she is giving an in-virtue-of explanations.

<sup>3</sup>One might claim that in-virtue-of explanations explain facts, not events, and so are not even potential counterexamples to the thesis that all explanations of events are casual. (This is in the spirit of something Lewis says ([1986], p. 223), though he discusses a different example and does not use the notion of an in-virtue-of explanation.) It is true that in the ‘distance from A to B’ example what is explained is the *fact* that A and B are 5 meters apart. But often an explanation that is presented using fact-talk can easily be converted into an explanation of a ‘corresponding’ event. The thesis that in-virtue-of explanations only explain facts must be the thesis

planations of events are causal is easily refuted. But the thesis that all explanations of events *other than in-virtue-of explanations* are causal remains an interesting and controversial thesis; it is the thesis I will defend.

(Since explanations of events will be my focus I should say something about the metaphysics of events. This much is uncontroversial: like ordinary material objects events have spatial and temporal locations. It makes sense to ask where and when an event happened. But after this all questions about the metaphysics of events are difficult. Among those questions are: how are material objects and events different (if indeed they are)? Can two or more events be spatially and temporally coincident? When do two event-names name the same event? It will not matter for my purposes what the answers to these difficult questions are, with one exception. I will presuppose a relatively liberal theory of events. Even if a region of space is completely empty from noon to 1pm on a given day I will say that there is an event consisting in that region's being empty during that stretch of time, an event about which we may ask, why did it occur?)

There are two things we need to figure out when we look at any example of an allegedly non-causal explanation: whether the example really is an example of an explanation, and whether it really is non-causal. How am I going to make those judgments? When judging whether a particular example is really an explanation I will not rely on any theory of explanation. I will just present the judgment that seems right to me. I do not think my judgments will be controversial.

On the other hand when judging whether a particular explanation is or is not causal I will use a theory. The theory I prefer is a kind of hybrid of two well-known theories of causal explanation. I am going to sketch the theory now and give some brief arguments in its favor. Then I will look at the examples.

I will work up to a statement of the theory in stages. A crude theory of causal explanation just says

---

that when an in-virtue-of explanation explains some fact F there is no corresponding explanation that explains the event that corresponds to F—because there *is* no event corresponding to F. (That was Lewis's view.) But in the theory of events I am assuming there is such an event as A and B's five-foot separation (it goes on as long as A and B are five feet apart) and such an event as the gas's pushing on the walls of its container with a pressure of 5 atmospheres.

(T1) A body of fact causally explains E iff it identifies a cause of E.

This theory is inadequate, for two reasons. First, as David Lewis notes (about an example I will discuss in more detail below), if some event E is uncaused then the fact that it is uncaused causally explains why it occurred (Lewis [1996]). The idea is that the information that some event was uncaused is the same kind of information as the information that some event was caused by (say) the alignment of the stars—it is information about an event’s causal history. And information about an event’s causal history is causal-explanatory information, even if the causal history is empty (the event is uncaused).

The second problem with (T1) is this: even if some event E has causes, a body of fact need not identify any of them in order to explain E. Suppose that a window breaks, and that Huey, Dewey, and Louie were the only three around who might have thrown a rock at it. The fact that Dewey did not throw but one of the other two did constitutes causal-explanatory information. But it does not identify the actual cause; it merely rules out one possible cause. Now maybe a *complete* causal explanation of the window’s breaking must say who threw the rock. But we should allow a body of fact to constitute a *partial* causal explanation even if it does not constitute a complete causal explanation. It is a partial causal explanation if it narrows down the list of possible causes (or possible causal histories) of the event being explained.<sup>4</sup> So we should move from (T1) to an improved theory:

---

<sup>4</sup>Sober ([1983]) requires causal explanations to identify an actual cause. He thinks that the more liberal theory I am describing ‘trivializes’ the notion of causal explanation: it makes pretty much any piece of information causal-explanatory information. His example, with slight modifications: consider a 10 meter tall flagpole. Sober claims that even the fact that the flagpole’s shadow is 15 meters long at 9am tells us ‘something about the cause’ of the flagpole’s being 10 meters high, namely ‘that it produced a [flagpole] that allowed the sun to cast the length shadow that it did’ (p. 203). It is true that the fact that the cause of the flagpole produced a flagpole that could cast a 15 meter shadow is a fact about the cause. But this fact does not narrow down the list of possible causes of the flagpole’s height so does not constitute a partial causal explanation on the current proposal.

Since Sober’s paper has come up, I should say in passing that his aim was to describe a class of explanations, equilibrium explanations, and argue that they are non-causal. But now the claim that equilibrium explanations are non-causal is widely

(T2) A body of fact completely causally explains E iff tells the complete story about what causes, if any, E has; a body of fact partially causally explains E iff it tells part of that story.<sup>5</sup>

But (T2) is still too weak—there are causal explanations that (T2) does not say are causal. The following example shows why. The water on my stove is currently 80 degrees Celsius. Why is it that hot? The only relevant cause of the water's temperature is the state of the knob that controls the burner (let us suppose). Now the knob's being turned half-way is what caused the water to be at 80 degrees. But even if we ignore influences on the water's temperature other than the knob the fact that the knob is turned half-way is not the largest body of explanatory information I could have. I would better understand why the water is at 80 degrees if I knew in addition the precise details of how the water's temperature depends on the knob's position: if I know the function  $f(r) = T$  that says what the water's temperature would be (in degrees Celsius) if the knob were turned  $r$  degrees from the off position. Knowing that the knob's being turned half-way caused the water to be at 80 degrees only gives me a tiny bit of information about this function: I only know that  $f(90) = 80$ , but not the value of  $f$  for any other argument.

Let me set out the moral of this story in more detail.<sup>6</sup> Knowing what events caused E at best puts you in a position to know what it would have taken for E not to have occurred.<sup>7</sup> But one gets additional explanatory information about E when one learns, in addition, what it would have taken for some specific alternative or rejected. Because it is discussed in great detail elsewhere (see (Strevens [2009]) I will not talk about this kind of example.

<sup>5</sup>Lewis proposes and defends a theory similar to this one in his paper 'Causal Explanation' (Lewis [1986]). As he notes his view is not that different from Railton's ([1981]).

<sup>6</sup>Something like this moral is pressed, in different contexts, by (Hitchcock [1996]) (he is more focused on causation than on explanation) and (Woodward [2003]).

<sup>7</sup>'At best' because there may be backup causes around ready to produce E should its actual causes fail. This is the problem of preemption for counterfactual theories of causation. It is an important problem but I will ignore it. Taking the problem into account explicitly would require me to say things in a more complicated way but would not change the claims I want to make in any important way.

range of alternatives of E to have occurred instead. And it would be absurd to say that this kind of information is not causal-explanatory information.

This brings us to the theory (really it is a theory-sketch) that I will use in this paper. It adds an extra clause to (T2) to allow the kind of information mentioned in the last example to count as causal-explanatory:

(T3) A body of fact partially causally explains E iff it is a body of fact about what causes, if any, E had; or if it is a body of fact about what it would have taken for some specific alternative or range of alternatives to E to have occurred instead.<sup>8</sup>

I should say that (T3) is a theory of the explanation of particular events. But scientists are also interested in ‘general’ explanations: they aim to explain not just why the earth orbits in an ellipse but why planets in general orbit in ellipses. (In fact science aims more at general explanations than particular ones.) But (T3) fits easily into a theory of general explanation: a general causal explanation says what the causal histories of instances of the event-type being explained have in common, or says something about what it would have taken for a given alternative type of event to have occurred instead that applies to many or most of the instances of the event-type being explained.<sup>9</sup>

Now (T3) says that information about an event’s causal history is causal-explanatory. But what is E’s causal history? It is the list of E’s causes, and the causes of its causes, and so on. But what are E’s causes? I am not going to describe

---

<sup>8</sup>Some philosophers (notably Hitchcock [1996], Woodward [2003], and Woodward and Hitchcock [2003]) think that examples like the ones I’ve discussed show that the objects of explanation, and the things that cause and are caused, are not events but ‘variables’. A variable is a kind of generalized event. An event must either occur or not occur in each possible world. We may say that *occurs* and *fails to occur* are the only two ‘values’ an event may have. Variables are like events except that there are more than just two values that they may take on. For example, *the water’s temperature (in Celsius)* is a variable: it can take any positive real number as a value. In this paper I will stick to the language of events. (There is a lot more to be said about what variables are; see for example Hall [2007].)

<sup>9</sup>This way of connecting particular and general causal explanations comes from (Lewis [1986]).



in detail a theory of causation. I am going to avoid as much as possible saying anything about causation and about any given event's causes that might be controversial (though I will have to do this at one point, when I talk about white dwarf stars in section 5.)

There are some interesting arguments for the existence of non-causal explanations that do depend on controversial claims about causation. For example, some philosophers think that there is no causation at the microphysical level. They think that causation is essentially a macroscopic-level phenomenon, or that the concept of causation only starts to get a grip on reality at a certain level of abstraction from exact microphysical details. But certainly physicists have found plenty of explanations of microphysical events. If there is no causation at the micro-level then these explanations are non-causal.<sup>10</sup> I am not going to say anything about this argument.

The notion of a partial causal explanation appears in (T3), and that notion will be important in what follows. I gave an example earlier of a partial causal explanation that tells just part of the story about E's causal history. It ruled out some hypotheses about what caused E without narrowing down the 'live' hypotheses to just one. That was a partial causal explanation that fit the kind of causal explanation described in (T2). For the sake of completeness I should mention that there can also be partial causal explanations that fit the kind of causal explanation that motivated the move from (T2) to (T3). The exact value of the function  $f$  (the one that gives the temperature of the water as a function of the angle of the knob) for each of its arguments is relatively complete causal-explanatory information. But information about the values of  $f$  that falls short of giving its exact value for each argument is still partial causal-explanatory information.

With these preliminaries out of the way we can now start discussing the examples.

---

<sup>10</sup>Field ([2003]) and Woodward ([2007]) express sympathy for this view about causation.

### 3 Explanations that Cite Causally Inert Entities

Graham Nerlich gives examples of what he says is a ‘geometrical, non-causal style of explanation’ ([1979], p. 74). Here is one of them. Imagine a cloud of dust particles moving through curved space. (Since we can visualize curved two-dimensional surfaces but not curved three-dimensional spaces, it may help to suppose that the cloud lives in a two-dimensional world.) The dust particles move inertially: no external forces act on them and they exert no forces on each other. So each particle travels along a ‘straightest’ line in curved space. But because space is curved the shape of the dust cloud keeps changing.

It seems that the curved geometry of space explains why the cloud’s shape changes. But space is causally inert.<sup>11</sup> It does not cause the cloud’s shape to change. And nothing else causes the cloud’s shape to change; for no forces are acting on the cloud. So here (Nerlich claims) we have a non-causal explanation.

But this example is not a non-causal explanation according to (T3), for reasons that have already come up. The cloud changes shape between  $t$  and  $t + \epsilon$ ; let the event E be this change. Does Nerlich’s explanation rule out any of E’s possible causal histories? Certainly. The explanation rules out all possible causal histories in which free particles do not travel on straightest lines but instead require external causes to do so. Nerlich says that E is uncaused, but it does not follow that his explanation is non-causal. For the proposition that E is uncaused rules out plenty of possible causal histories.

Nerlich says that this explanation is an instance of a ‘geometrical’ style of explanation. But the explanation does not only tell us about the geometry of space. It also tells us something about the dynamical laws governing the behavior of matter in that space. (For one thing, it tells us that the law of inertia is true.) And the dynamical laws help determine the causal facts. No surprise, then, that there is information about E’s causal history in this explanation after all.

What about the part of Nerlich’s explanation that just says something about the geometry of space? Is that part of the explanation, by itself, partially explana-

---

<sup>11</sup>At least, this is true in the physics we are imagining. Things might be different in the general theory of relativity.

tory? It need not be. 'P partially explains E' does not follow from 'P and Q explain E'. And even if a proposition about the geometry of space does partially explain E that proposition would not do so non-causally. The proposition that space is curved in such-and-such a way does rule out some of E's causal histories. In hyperbolic geometry there are no rectangles. So the proposition that space is hyperbolic rules out causal histories in which the dust cloud once had a rectangular configuration. The moral here goes back to the difference between (T1) on the one hand and (T2) and (T3) on the other. The fact that space has such-and-such a geometry does not seem to be a cause of any of the cloud's behavior; but a proposition can rule out causal histories without in any intuitive sense citing a cause.

Another alleged example of geometrical, non-causal explanation, this one due to Peter Lipton, does not work for similar reasons:

[S]uppose that a bunch of sticks are thrown into the air with a lot of spin, so that they separate and tumble about as they fall. Now freeze the scene at a moment during the sticks' descent. Why are appreciably more of them near the horizontal axis than near the vertical, rather than in more or less equal numbers near each orientation as one might have expected? The answer, roughly speaking, is that there are many more ways for a stick to be near the horizontal than near the vertical [...] there are two horizontal dimensions but only one vertical one. (Lipton [2004], pp. 31-2)

Lipton says that this is a non-causal explanation because geometrical facts cannot be causes. But we have just seen that (T3) does not say that geometrical facts cannot causally explain. And, as with Nerlich's example, the explanation does not rely just on geometrical facts; dynamical facts also play a crucial role in Lipton's explanation.

To see this suppose that sticks 'strive' to orient themselves horizontally. Then most of the sticks would be nearly horizontal even if the number of horizontal dimensions and vertical dimensions were the same (this could happen if the sticks lived in a two-dimensional space). Clearly if that is how the dynamics of the sticks works then the geometrical facts are explanatorily irrelevant. To the extent that the

geometrical facts appear relevant it is because we assume that the dynamics governing the stick's behavior works differently. We assume that the dynamics does not prefer one spatial direction over another.<sup>12</sup>

Here is a third example that is often said to be an example of a geometrical, non-causal explanation (one philosopher who says this is Colyvan ([2001], p. 50). Special relativity says that 'moving sticks shrink'. The geometry of relativistic (Minkowski) spacetime (supposedly) explains why.

Let's get a closer look at the phenomenon and its explanation. The spacetime diagram in figure 1 gives the setting. The two bold lines are the worldlines of the end of a meter stick. (The worldlines are straight; the meter stick is in unaccelerated motion.) One observer has as his worldline L1; he is always right next to one end of the meter stick. Another observer has as his worldline L2; he moves past the meter stick. In the L1 frame of reference the stick's length is one meter. But in the L2 frame of reference the stick's length is less than one meter. The phenomenon to be explained, then, is: why is the stick shorter in the L2 frame of reference than it is in the L1 frame of reference?<sup>13</sup>

The geometrical explanation goes like this. The fundamental geometrical notion in relativistic spacetime is the spacetime interval between two points. This notion is used to define the spatial distance between two spacetime points in a frame of reference. This definition says that the length of the stick in the L1 frame of reference can be calculated like this: find two spacetime points that (i) lie on the ends of the stick, and (ii) are simultaneous in the L1 frame of reference. Then the length of the stick is the square-root of the spacetime interval between those two points. In the diagram points A and B are simultaneous in the L1 frame, so writing  $I(A, B)$  for the interval between A and B, the length of the stick in the L1 frame is  $\sqrt{I(A, B)}$ .


Similarly, the length of the stick in the L2 frame of reference is the spatial

---

<sup>12</sup>Lange makes a similar point about this example (Lange [2009], p. 420).

<sup>13</sup>There is something else that one might want explained: why, in a single frame of reference, a stick moving in that frame is shorter than a stick at rest in that frame (where the two sticks are the same length when placed side by side at rest). But the geometrical explanation of this fact is in all essential respects the same as the one I will look at.

Figure 1: Diagram of Length Contraction



distance between two spacetime points that lie on the ends of the stick and are simultaneous in that frame of reference. A and C are a pair of points like that, so the length in the L2 frame of reference is  $\sqrt{I(A,C)}$ .

To complete the explanation we just need three more facts from relativistic geometry: first, that line segments AB and BC form a ‘relativistic right angle’, and so triangle ABC is a ‘relativistic right triangle’; second, that the analogue of the Pythagorean theorem holds for relativistic right triangles, with the interval playing the role of the squared-distance:  $I(A,B) + I(B,C) = I(A,C)$ ; third, and finally, that  $I(B,C)$  is a *negative number* (the spacetime interval between any pair of points that both lie on the worldline of some massive body is negative). It follows from these three claims that the stick’s length in the L2 frame is less than its length in the L1 frame.

Even if this is a non-causal explanation it does not show that there are non-causal explanations of the kind we are investigating. For this explanation is quite obviously an explanation of the wrong kind. It is an in-virtue-of explanation. The explanation cites the geometrical facts that ground the fact that the stick is shorter in the L2 frame than in the L1 frame.

(What would an explanation of length contraction that was of the ‘right’ kind look like? It is easy to imagine a causal explanation. Suppose that the stick ended up following its actual path through spacetime because someone pushed on it. If they had pushed harder then the left end of the stick would have had L2 as its worldline instead of L1. If that had happened the stick would have been *longer* in the L2 frame than in the L1 frame. So the softness of the push (causally) explains why the stick is longer in the L1 frame than in the L2 frame.)

Now there has been a lot of debate recently over the role of spacetime geometry in explanations of length contraction. Brown and Pooley ([2006]; see also Brown [2005]) claim that geometrical facts cannot explain while Balashov and Janssen ([2003]) claim that they can. But this debate is not relevant to what I have been discussing. First, Brown and Pooley *agree* that the explanation that I have been discussing really is an explanation: about it and others like it they write ‘In our opinion these constitute perfectly acceptable explanations (perhaps the only acceptable explanations) of the explananda in question’ (p. 78).<sup>14</sup> Their claim (as I understand it) is that there are other explanatory requests, which can also naturally be expressed by the words ‘please explain relativistic length contraction’, which cannot be answered by citing facts about the geometry of spacetime. And nothing I have said is incompatible with this claim. In the end the debate between Brown and Pooley on one hand and Balashov and Janssen on the other is a debate about ‘the direction of explanation’ between the dynamical laws and the geometry of spacetime: does the geometry explain the fact that the laws are Lorentz invariant, or the other way around? But that debate is not a debate about the status of any explanation of a particular event.

#### **4 Explanations that Merely Cite Laws, I**

Another place people have looked for non-causal explanations is in the class of explanations that cite laws without explicitly saying anything about what caused the

---

<sup>14</sup>However, they probably disagree with me that this explanation is an in-virtue-of explanation. For they think that the behavior of matter grounds facts about the geometry of spacetime; so facts about the geometry of spacetime cannot also ground any facts about the behavior of matter.

target event: explanations that ‘merely’ cite laws. Examples like this seem to favor Hempel’s Deductive-Nomological theory of explanation (Hempel [1965]) over the thesis that all explanations are causal. Here is one example of this kind: ‘That meteor is traveling slower than light because no material thing can travel faster than (or at) the speed of light’. This is certainly an explanation.

But even if an explanation cites a law while saying nothing explicit about E’s causal history it may still convey information that (T3) counts as causal-explanatory. The meteor explanation cites no cause of the meteor’s subluminal velocity, but does rule out causal histories in which meteors left alone always accelerate to a superluminal speed and something special happened to prevent this one from doing so.<sup>15</sup> Similarly, ‘that bird is black because it is a raven and it is a law that all ravens are black’ (if it is a genuine explanation) rules out causal histories of that raven’s being black in which raven genetics do not require ravens to be black (it was not born black but was later painted black by a mischievous epistemologist).

Now James Woodward denies that ‘That meteor is traveling slower than light because no material thing can travel faster than (or at) the speed of light’ is a causal explanation. That is because he does not think that ruling out causal histories is sufficient for being causal-explanatory. He says instead that a causal explanation of E must identify some prior conditions (causes of E) and say that if those conditions had been changed then E would not have occurred. The meteor explanation does not do that. Here is how he puts it in a discussion of a similar example: ‘explanations that appeal to [the relativistic prohibition] fail to identify the conditions on which the (sub)luminal velocity of the particle [or meteor] counterfactually depends. To the extent that saying what an outcome depends on is at the heart of successful [causal] explanation, [the relativistic prohibition] fails to be [causally-] explanatory’ (Woodward [2003], p. 209).<sup>16</sup>

---

<sup>15</sup>Strictly speaking it is controversial whether special relativity says that no material thing can travel faster than the speed of light. But it certainly does say that *meteors* (which are made of ordinary matter—protons, electrons, neutrons) cannot travel faster than the speed of light.

<sup>16</sup>It is actually not clear to me from the text whether Woodward merely claims that the meteor explanation is not causal, or whether he makes the stronger claim that it is not an explanation at all. What he actually writes is that it ‘fails to be

As I have said, this seems to me an unnecessarily strong requirement to put on causal explanation. If I am wondering why the window broke and I learn that Huey did not throw a rock at the window then I have learned *a little bit* at least about why the window broke, because (given my background knowledge) I now know that it was either Dewey or Louie. True, the information I learn does not identify a condition on which the breaking depends; but it gets me closer to identifying such a condition.

I have said that the meteor explanation is causal because it rules out causal histories incompatible with the cited law. Here is another way to put this reason for regarding the meteor explanation as causal, one that may make it more appealing. (T3) says that information is causal-explanatory if it says something about what it would have taken for the meteor to move faster than the speed of light. While the meteor explanation does not give a complete story about what it would have taken for that to happen it does say something: faster-than-light motion would have required different physical laws.

Now if one wants to know why the window broke then the fact that Huey did not throw is only a little explanatory; but if one wants to know why the meteor is moving at less than the speed of light the fact that no material thing can travel faster than light is (to put it mildly) more than a little explanatory. What is the difference? I think that Woodward's own insights into the nature of causal explanation show why the citation of the relativistic prohibition is such a good causal explanation. The relativistic prohibition entails that no physically possible interventions on any prior conditions would result in the meteor's moving faster than light. The prohibition entails that the meteor's subluminal velocity is counterfactually *independent* of all prior conditions (so long as we stay within the realm of physical possibility).  

---

explanatory', without the 'causal' qualifier. But he also says earlier in the book that this qualifier is implicit even when it does not appear (p. 187). Still, not everything he says makes sense to me if we assume that by 'explanation' he always means 'causal explanation'. For he admits that some think it counterintuitive to say that citing the relativistic prohibition is not explanatory (p. 209). Certainly it is counterintuitive to say that the example fails to be explanatory *at all*. But who thinks it is counterintuitive to say that it is not a *causal* explanation, when it seems at least initially like such a good candidate for a non-causal explanation?



Now Woodward says that a causal explanation must identify prior conditions on which the target event depends. But in this case *there are none* (that are physically possible). The information that there are none certainly looks like the same kind of information as information that identifies conditions on which the target event depends. In the limiting case where *there are no such conditions to identify* the most you can say is that there are no such conditions. This information should therefore count as causal-explanatory information. And it should count as a particularly excellent kind of causal-explanatory information. An explanation of some event that identifies one or a few of its causes is incomplete; there are many other causes that a more complete explanation might identify. But an explanation that, like the meteor explanation, says that the target event was (as it were) unpreventable says something about all of its causes: change any of them in any relevant way, and the event still happens. (In the case of the meteor the relevant ways are the physically possible ways of changing prior conditions that are compatible with that meteor's being there at all.)

## 5 Stellar Collapse

I have been talking about an explanation that merely cites a law. There are others like it that have been offered as examples of non-causal explanations. What I want to say about them is similar to what I have said about the meteor example. But there is one example that I want to discuss in depth, because it is so well-known and because I think the previous discussions of it have been confused. David Lewis made the example famous; here is how he puts it:<sup>17</sup>

A star has been collapsing, but the collapse stops. Why? Because it's gone as far as it can go. Any more collapsed state would violate the Pauli Exclusion Principle. It's not that anything caused it to stop—there was no countervailing pressure, or anything like that. There was nothing to keep it out of a more collapsed state. Rather, there just was no such state for it to get into. The state-space of physical possibilities

---

<sup>17</sup>The example appears in (Lewis [1986], p. 222); Lewis attributes it to Peter Railton.

gave out.

Let E be the event *the cessation of the collapse*. This (purported) explanation says that (i) E was not caused (or at least that it had no causes ‘except for all the causes of the collapse which [were] a precondition of the stopping’) and that (ii) when E occurs the star had reached the boundary of its state space. Lewis latches on to the first part of the explanation and simply says that the explanation thereby rules out possible causal histories of E (namely, histories in which it is caused).

There are a lot of puzzling things about this example. The first problem with it that we should attend to is that the causal claims made in the example are wrong. It is just not true that nothing caused the star to stop collapsing, that ‘there was no countervailing pressure or anything like that’. One physics textbook says this about white dwarf stars: ‘What supports [the matter making up the star] against gravity? The short answer is: the pressure of a degenerate electron gas’ (Baierlein [1999]. p. 194). So there *was* a countervailing pressure. (A collapsing star can be modeled very accurately as a cold electron gas. Even though a collapsing star is very hot by ordinary standards for thermodynamic purposes it counts as being close to absolute zero.) Of course, as another textbook warns, ‘degeneracy pressure has absolutely nothing to do with electrostatic repulsion between the electrons [...] it arises purely by virtue of the exclusion principle’ (Schroeder [2000], p. 275). That is: the degeneracy pressure of a cold (degenerate) electron gas does not come from the electrons’ mutual negative charge pushing them apart—it does not come from one of the ‘usual’ forces. This might lead one to say that electron degeneracy pressure is not ‘real’ pressure, or that it is some kind of fictional pressure, or even that really there is no such thing as electron degeneracy pressure at all. If these claims are true then perhaps the explanation *does* have its causal facts right.

But these claims are not true. If the only theory in which the term ‘pressure’ appeared was the (classical) kinetic theory of gasses then it might be plausible to define ‘pressure’ so that the gas’s ‘true’ pressure is due only the electrons’ mutual electrostatic repulsion. If our theories were limited in that way it might also be plausible to define ‘temperature’ as ‘mean molecular kinetic energy’. But these terms do not just appear in the kinetic theory of gases. They appear in statistical mechanics, a theory that applies to lots of physical systems besides classical gasses.

The successes of that theory show that pressure, like temperature, is a much more abstract quantity than the kinetic theory might lead one to believe. (Idealized) statistical mechanical systems in which the mean molecular kinetic energy is always very small, like ideal paramagnets, can still have very high temperatures. What then is temperature? A system's temperature is defined to be  $\partial U / \partial S$ , the partial derivative of (internal) energy with respect to entropy (at fixed external parameters). This quantity coincides in a certain respect with mean molecular kinetic energy in some gases, but not in other systems. The statistical mechanical definition of 'pressure' is similarly abstract: it is defined to be  $-\partial U / \partial V$ , the partial derivative of the internal energy with respect to volume (at fixed entropy). These definitions reflect the fact that in statistical mechanics facts about a system's temperature and pressure are not facts about the system's 'occurrent' properties: a system's temperature is not a matter of how much energy it contains and a system's pressure is not a matter of how hard the system's constituents are banging on the container. Instead these quantities are *dispositional*: a system's temperature is a matter of how disposed it is to *transfer* energy to another body, and a system's pressure is a matter (roughly speaking) of how disposed it is to 'transfer' volume to another body. And just as things other than the kinetic energy of a system's constituent molecules can contribute to its temperature, things other than the repulsive forces between its constituent molecules—things like the consequences of the Pauli Exclusion Principle—can contribute to its pressure.<sup>18</sup> (I will briefly explain how they do so presently.)

So the Railton / Lewis 'explanation' is wrong to say that E is uncaused. And since explanations have to be true (or at least approximately true) it follows that this 'explanation' is no explanation at all.

Remember, there were two parts to the 'explanation': (i) E is uncaused, and (ii) when E occurs there are no more collapsed states for the star to get in to. We have seen that (i) is not true. A natural thought to have is to remove (i) from the explanation and just stick with (ii). Why not just say that the star stopped collapsing because there were no more collapsed states available? If that explains why the star stopped collapsing then we have an explanation that does not (explicitly) say that the cessation was uncaused, and it is no longer obvious that Lewis's response to the

---

<sup>18</sup>Baierlein ([1999], p. 349) makes these points with respect to temperature.

example is a good one.

But it turns out that (ii), like (i), is false.

The Pauli Exclusion Principle says that no two identical fermions (electrons are fermions) can occupy the same quantum state. Now if one is thinking classically then one expects that if a gas's temperature is very close to absolute zero then the gas's constituent particles will all be in very low energy states—none of them will be moving very fast and almost all of them will be moving very slowly. But that is not what happens to a cold electron gas when we take quantum mechanics into account. There are relatively few lowest energy states an electron can be in, and the Exclusion Principle says that each of those states can have only one electron in it. So even if the gas's temperature is very close to zero many of its constituent electrons will be in relatively high energy states. In this way the Exclusion Principle plays a role in determining the internal energy of an electron gas and thereby (via the equation  $P = -\partial U / \partial V$ ) its pressure.

What the Exclusion Principle excludes are states of the star in which more than one of the electrons are in the same quantum state. But the star can occupy a smaller volume without that happening. Decreasing the volume of the star does change how many single-electron states there are with a given energy level. But there are infinitely many energy levels and so infinitely many states available. So even if the star occupied a smaller volume there is nothing stopping each electron from getting into an available state: figuratively speaking the first electron occupies the lowest energy state and later electrons occupy higher and higher energy states in turn until all the electrons are in some state or other.<sup>19</sup> No matter how small the star's radius the electrons never run out of states because there are infinitely many of them.

To sum up: the (revised) Railton / Lewis 'explanation' is wrong—there *are* more collapsed states for the star to get in to. Once again, since we do not have an explanation of E here we do not have a non-causal explanation.

So why does the star stop collapsing? Not because a state with a smaller radius is physically impossible, but because the star has reached the radius at which

---

<sup>19</sup>Of course strictly speaking it makes no sense in quantum mechanics to ask which electron occupies which quantum state.

the outward-directed pressure in the star exactly balances the inward-directed gravitational forces. This is a paradigmatically causal explanation.

It is true that there are states the star cannot get in to: states in which all the constituent electrons are in the (or a) lowest energy state. But the role this fact plays in the explanation of E is different from the role the Pauli Exclusion Principle is represented as playing by the ‘explanation’ Lewis and Railton give. The way they tell the story the star looked like it was ‘heading for’ one of the prohibited states. Then the star ‘stopped just short’ of entering one of the prohibited states, and stopped because that state was prohibited. Nothing like this is in fact true. I do not know how to make sense of the idea that the star was at some point ‘heading for’ a state in which two electrons occupied the same quantum state.

## **6 Explanations that Merely Cite Laws, II**

The Lewis/Railton ‘explanation’ is still of interest to us even though it contains so many false claims. There is an abstract form of explanation that it is trying to exemplify, a form that Railton suspected was non-causal. Even if the true explanation of E fails to exemplify that form, other actual or possible explanations (of other real or hypothetical events) might. So we should try to figure out what this form is and whether explanations that instantiate it are non-causal.

It is not hard to figure out what that form is. The state-space of a physical system is the set of all the instantaneous states that it is physically possible for systems of that kind to get in to. State spaces (typically) have some topological structure (in fact they typically also have some kind of geometrical structure), and a state space may have a boundary or ‘edge’. Abstractly speaking, the kind of situation we want to focus on is one in which a physical system is ‘heading for’ the boundary of its state space (better: the point in state space that represents the system’s state is heading for the boundary). It reaches the boundary and then stops: it remains in a state S right on the boundary, rather than going into a state on the other side of the boundary.<sup>20</sup> The question is: why does the system end up in state

---

<sup>20</sup>I speak here as if there are states on the other side but it is physically impossible to reach them. Railton speaks as if there *are* no states on the other side. I do not think much hangs on which way of talking we adopt in this context. I also speak

S?

Imagine a particular hypothetical system that ends up in a particular hypothetical boundary state S. Let F be the event: the system's state stops changing and ends up in S. Now Lewis's idea (adapted to the present context) is that (a) F is uncaused and (b) that its being uncaused causally explains why it happened. I have already accepted that (b) is true if (a) is true. But what about (a)—is it true that F is uncaused?

The example has been presented at too abstract a level to figure this out. We have been told some things about the system's state space (namely, that it has a boundary), but nothing about the dynamical laws for the system. And what the dynamical laws are surely makes a difference to what the causal facts are.

Suppose that it is a law of nature that everything is confined to the inside and boundary of a giant sphere. (There is space outside the sphere, but nothing is allowed to enter that region of space.) And suppose one particular body is heading toward the boundary but then slows down and stops just at the boundary. Is its stopping uncaused? Not if the dynamics works like this: near the center of the sphere Newton's laws are approximately correct. Bodies with no other bodies nearby are not subject to any strong forces and so move on nearly straight lines at nearly constant speed. But near the boundary things are different. When something gets close to the boundary it experiences a kind of 'friction' with space itself, and slows down. The magnitude of the friction depends on the body's speed and its proximity to the boundary in such a way that a body heading toward the edge of the sphere always stops before crossing the boundary. Clearly when the body stops at the boundary the stopping is caused by the friction.

This example shows that something's stopping at the boundary of its state space need not be uncaused. (Even if one thinks that the scenario described is impossible—one might balk at the notion of friction with space itself—surely there are possible scenarios that are structurally isomorphic to it.) Of course this example is no problem for the thesis that all explanation of events is causal. It is only meant 

---

 here as if the laws permit that system to occupy states on the boundary, even though it is more natural for the laws to permit the system only to occupy states within the boundary (since differential equations are defined on open sets).

to show that one should not infer that the stopping was uncaused just because it was a stopping at the boundary of state space.

Why might this inference have seemed good? One might think that if the system *cannot* cross the boundary then no cause is needed to prevent the crossing. But even if this is true it does not follow from the fact that no cause is needed that no cause exists.

Lying behind this line of thought might be a conception of the way dynamical laws must deal with boundaries in state space. But that conception is not correct in all cases. The conception goes like this: the dynamical laws, if left to operate ‘on their own’, would compel the system to cross the boundary. So the law asserting the existence of the boundary needs to ‘step in’ and stop a violation that would otherwise occur. The dynamical laws are ‘suspended’ in favor of the law describing the boundaries of state space. And events that occur when the dynamical laws are suspended are uncaused.

But dynamical laws and laws about the boundaries of state space do not need to be independent of each other in this way. The fanciful example above shows this. For all that has been said about the example the law asserting that nothing can leave the giant sphere need not be an independent fundamental law. It might just be a derivative law that follows from (and is explained by) the dynamical laws (and perhaps the initial conditions—the fact that the universe started out with everything inside the sphere).

(There is a third possibility too. Maybe the law asserting that nothing can leave the sphere is more fundamental than the dynamical laws, and explains why the dynamical laws prevent anything from leaving the sphere.<sup>21</sup>)

---

<sup>21</sup>Which of these possibilities describes the relationship between the Pauli Exclusion Principle and the dynamical laws in quantum mechanics (Schrödinger’s equation)? I am not sure. It is certainly either the second or the third, not the first. Schrödinger’s equation does not allow a quantum state that satisfies the Pauli Exclusion Principle to evolve into one that does not. The Pauli Exclusion Principle does not need to step in and stop a violation that would otherwise occur. Whether to think we have here an instance of the second or the third possibility might depend on the relationship between the Pauli Exclusion Principle and the requirement of permutation invariance: see (Redhead [1992]) and (Huggett [19992]) for discus-

I have been talking about what explains why a body stopped at the boundary of the giant sphere. Depending on the dynamical laws the stopping may have been caused, or it may have been uncaused. In either case there is a causal explanation to be given. But Railton's original idea, I think, was that there is another explanation one might give, distinct from either of these causal explanations. What is this other explanation? I can think of two candidates.

The first explanation just cites the law giving the boundaries of state space and remains silent about the dynamics:

- (1) The body failed to cross the boundary because it is a law that nothing crosses the boundary.

This explanation does not make the mistake that the Lewis / Railton explanation made of citing the law describing the boundary and then *adding* the claim that the stopping was uncaused. It is because this explanation remains silent about whether the stopping was caused or uncaused that it seems like a non-causal explanation.

I do not think this this example is that hard to handle. It is relevantly similar to the meteor explanation, and is causal for the same reasons. It tells us that the stopping was unpreventable by any physically possible means (that are consistent with the body's moving toward the boundary in the first place). This kind of explanation, I have already argued, is causal.

I said that I can think of two candidate explanations of the stopping that at first glance look like they can serve Railton's purpose. The second is more detailed than the first and draws on ideas I have already discussed:

- (2) The body failed to cross the boundary because it is a law that nothing crosses the boundary; and this law is independent of the dynamical laws, and 'trumps' them when they conflict with it.

In this example the boundary law and the dynamical law are related in the first way I discussed above.

Suppose (2) is true (in the hypothetical world we are imagining). Then it certainly looks explanatory. Like (1) this explanation says nothing explicit about the

---



cause of the stopping. But (2) goes beyond (1) and it does so in a way that (one might think) suggests that it is non-causal. What (2) appears to do is identify something on which the stopping counterfactually depends: the law about the boundary. Because this law is independent of the dynamical laws, it is plausible to think that if the law had not been true then the body would not have stopped. For if the law had not been true the dynamical laws still would have been laws (by independence), and since (again by independence) the dynamical laws do not respect the boundaries given by the boundary law, they would have compelled the body to cross the boundary.

That the stopping counterfactually depends on the law suggests that the law explains the stopping; but the law is not an event, so is not a cause of the stopping, so does not causally explain the stopping. This second example, then, is different from the meteor example, because a similar counterfactual is not true in that example.

I suspect that this example is the best version of the kind of example Railton was trying to give. Is it an example of an explanation and if so is it really non-causal?

I have my reasons for doubting that (2) is an example of an explanation. It uses hypothetical physics, in fact a bare sketch of some hypothetical physics. Have we really described a possible way for the world to be? Is it really possible for boundary laws to be independent of dynamical laws and to trump them when they conflict? If not then (2) is not an explanation. (The good thing about the stellar collapse example, despite its flaws, is that it used (or tried to use) real physics.)

If (2) is an explanation, is it non-causal? One might wonder whether the claims about counterfactuals that give the example its force are true. Counterfactuals of the form ‘If a certain law of nature had been false...’ are difficult to evaluate. If the relevant counterfactuals are false then the reason given for thinking that (2) is non-causal goes away.

But let me set these worries about the example aside. Supposing that (2) is an explanation, I do not think it is non-causal, because I think a version of Lewis’s original reply works against it. The details about the relationship between the boundary law and the dynamical law that the explanation includes entail (for reasons I gave above) that the stopping at the boundary was uncaused. Since this explanation en-

tails something about the causal history of the stopping it is a causal explanation. And I think the kinds of things I said about the meteor example and about (1) apply here too: whatever the relationship between the laws, the boundary law entails that the stopping was unpreventable.

If there is one objection to my view it might go like this: the features of (2) in virtue of which it qualifies as a causal explanation are not the features of the explanation in virtue of which it is explanatory. It qualifies as explanatory, the objection goes, in virtue of the counterfactual dependence of the target event on the law.

I feel the force of this objection, but I ultimately reject it. My reservation is that I do not know how to tease apart the features of (2) that (the objection alleges) are the source of its explanatory power from the features of (2) that qualify it as causal. In some cases it is easy to do this: if an explanation is a conjunction, if it has the form 'X because A and B', it might be that only B says anything about X's causes. Then we could consider A alone and ask ourselves whether A by itself explains X. If it does we have found a non-causal explanation. But we cannot do this in the present case. I do not see how to remove from (2) the claim that the stopping was uncaused and unpreventable while leaving in the claim that the stopping was required by law. So I stick to my original conclusion: this explanation is causal.

## 7 A Final Example

The examples from the last three sections were examples that merely cited laws without saying anything explicit about what did or did not cause the events being explained. These laws, though physically necessary, are metaphysically contingent. There are other examples of allegedly non-causal explanation that cite *metaphysical* necessities, like mathematical truths. These examples do not raise any new issues but for the sake of completeness it is worth discussing one explicitly.

Colyvan ([2001], p. 49) proposes this example: right now there are two points  $p$  and  $q$  on the earth that (i) are antipodal (opposite one another), and (ii) have the same temperature and atmospheric pressure. Let E be the event: there currently being a pair of points like this. Why does E occur? The Borsuk-Ulam theorem

says that if  $f$  is any continuous function from a sphere (or any surface topologically equivalent to a sphere) to ordered pairs of real numbers then there are antipodal points  $x, y$  that get mapped by  $f$  to the same pair of numbers. The function from the surface of the earth to the temperature and pressure of the atmosphere at that point is such a function. So E *has* to occur; that is why it occurs.

One way to respond to this example is to deny that there is such an event as E. Why deny this? E is equivalent to the event: *this* pair of antipodal points having identical values, or *that* pair of antipodal points having identical values, or...; and so E is ‘too disjunctive’ to be a genuine event. I admit that this strategy is appealing. Supposing that E is an event, does it even have a causal history? While I can imagine events that could cause the north and south poles to have identical temperatures and pressures, I find it hard to imagine an event that could cause it to be that *some* pair of antipodal points is like this that is not also cause of this being true of one particular pair. But then maybe that means that E is not an event at all.

But I want to grant that there is such an event as E to be explained. Even still I do not think we have a non-causal explanation here. True, the mathematical theorem does not rule out any causal histories of E. But we now know how to handle this kind of case. The mathematical theorem tells us that E was unpreventable: nothing, not even changes to prior conditions that broke the laws of nature (if changes of that sort even make sense), could have produced any alternative to E. That kind of information is causal-explanatory information.

## 8 Conclusion

I have discussed many of the standard examples believers in the existence of non-causal explanations (of events) have offered. I do not think any of them work. They all rely on overly-restrictive views about what a causal explanation must do. A causal explanation does not need to identify a cause of the target event. It does not even need to imply that the target event was caused. And (I have argued) the fact that some event was independent of prior circumstances, that it was unpreventable, can causally explain that event.

## Acknowledgements

Thanks to everyone at MIT, especially Steve Yablo and Agustin Rayo; to the philosophers at Columbia University, especially David Albert; to Brad Weslake; and to three anonymous referees for detailed comments.

## References

Baierlein, R. [1991]: *Thermal Physics*, New York: Cambridge University Press.

Balashov, Y. and Janssen, M. [2003]: ‘Presentism and Relativity’, *The British Journal for the Philosophy of Science*, **54**, pp. 327-46.

Bennett, J. [1988]: *Events and their Names*, Indianapolis: Hackett.

Brown, H. R. [2005]: *Physical Relativity: Space-time Structure from a Dynamical Perspective*, New York: Oxford University Press.

Brown, H. R. and Pooley, O. [2006]: ‘Minkowski space-time: a glorious non-entity’, in D. Dieks (ed), 2006, *The Ontology of Spacetime*, Amsterdam: Elsevier, pp. 67-89.

Colyvan, M. [2001]: *The Indispensability of Mathematics*, New York: Oxford University Press.

Field, H. [2003]: ‘Causation in a Physical World’, in M. J. Loux and D. W. Zimmerman (eds), 2003, *The Oxford Handbook of Metaphysics*, New York: Oxford University Press, pp. 435-60.

Hall, N. [2007]: ‘Structural Equations and Causation’, *Philosophical Studies*, **132**, pp. 109-136.

Hempel, C. G. [1965]: *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, New York: The Free Press.

Hitchcock, C. [1996]: 'Farewell to Binary Causation', *Canadian Journal of Philosophy*, **26**, pp. 267-282.

Huggett, N. [1999]: 'On the Significance of Permutation Symmetry', *The British Journal for the Philosophy of Science*, **50**, pp. 325-47.

Lange, M. [2009]: 'A Tale of Two Vectors', *Dialectica*, **63**, pp. 397-431.

Lewis, D. [1996]: 'Causal Explanation', in his *Philosophical Papers Volume 2*, New York: Oxford University Press, pp. 214-40.

Lipton, P. [2004]: *Inference to the Best Explanation*, 2nd edition, New York: Routledge.

Nerlich, G. [1979]: 'What Can Geometry Explain?', *The British Journal for the Philosophy of Science*, **30**, pp. 69-83.

Railton, P. [1981]: 'Probability, Explanation, and Information', *Synthese*, **48**, pp. 233-56.

Redhead, M. and Teller, P. [1992]: 'Particle Labels and the Theory of Indistinguishable Particles in Quantum Mechanics', *The British Journal for the Philosophy of Science*, **43**, pp. 201-18.

Rosen, G. [2010]: 'Metaphysical Dependence: Grounding and Reduction', in B. Hale and A. Hoffman (eds), 2010, *Modality: Metaphysics, Logic, and Epistemology*, New York: Oxford University Press, pp. 109-35.

Salmon, N. [1984]: *Scientific Explanation and the Causal Structure of the World*,

Princeton: Princeton University Press.

Schroeder, D. V. [2000]: *An Introduction to Thermal Physics*, New York: Addison Wesley Longman.

Sober, E. [1983]: 'Equilibrium Explanation', *Philosophical Studies*, **43**, 201-10.

Strevens, M. [2009]: *Depth: An Account of Scientific Explanation*, Cambridge, MA: Harvard University Press.

Woodward, J. [2003]: *Making Things Happen: A Theory of Causal Explanation*, New York: Oxford University Press.

Woodward, J. [2007]: 'Causation with a Human Face', in H. Price and R. Corry (eds), 2007, *Causation, Physics, and the Constitution of Reality*, New York: Oxford University Press, pp. 66-105.

Woodward, J and Hitchcock, C. [2003]: 'Explanatory Generalizations, Part I: A Counterfactual Account', *Nous*, **37**, pp. 1-24.