

Causal Pluralism

Peter Godfrey-Smith

Philosophy Department
Harvard University

October 2007

To appear in the *Oxford Handbook of Causation*,
edited by Helen Beebe, Chris Hitchcock, and Peter Menzies.

1. Introduction

2. Methodological and Other Preliminaries

3. Two Concepts of Cause?

4. The Amiable Jumble

5. Anscombe, Minimalism, and Family Resemblance

6. Causation as an "Essentially Contested" Concept?

1. Introduction

Causal pluralism is the view that causation is not a single kind of relation or connection between things in the world. Instead, the apparently simple and univocal term "cause" is seen as masking an underlying diversity.

Assessing such a claim requires making sense of a difficult counting operation. How do we tell whether a theory of causation is identifying causation with a "single" kind of connection? In practice, there tends not to be much disagreement about how to do the counting, because most philosophical work on causation has sought a view with an obvious kind of unity. The literature often works with a standard range of *candidate* connections that seem to have an important link to the idea of causation. These include natural laws, *de facto* regularities, counterfactual dependence, probabilistic dependence,

and some others. It has been common in philosophy to pick one of these and try to make sense of causation entirely in those terms. The candidate chosen is seen as either fundamental to our thinking about the world, a fundamental ingredient of the world itself, or both. Against that background, it is an unorthodox move to say that all such unified treatments of causation are mistaken; causation is, with respect to the ingredients recognized as distinct options in philosophical practice, irreducibly plural or diverse.

Appropriately, causal pluralism comes in several kinds. The main focus of this survey is a collection of recent treatments that directly oppose the kind of unity seen in traditional analyses. The simplest kind of pluralism treats our ordinary talk of causation as shifting between two distinct "concepts" of cause (section 3). This view posits an ambiguity that could be resolved if each concept was given a different name. Another suggestion is that we have a single concept of causation, but one whose use is guided by what Brian Skyrms called an "amiable jumble" of criteria that sometimes work together and sometimes pull apart (section 4).

Those two options use traditional philosophical raw materials to make sense of causation (regularities, counterfactuals, etc.) but put them to work in unorthodox ways. A more unusual style of analysis is seen in a proposal by Anscombe. Perhaps "cause" is a generic term whose application to the world is parasitic on the semantic properties of a range of terms (such as "burn" and "scrape") that pick out specific kinds of physical connection. Anscombe's proposal raises the possibility of a kind of "minimalism" about causation, and also a "family resemblance" view that is closely related (section 5).

After surveying these options, I suggest that an understanding of the workings of our concept of causation may require that we treat it as something like what W. B. Gallie (1956) called an "essentially contested" concept. These are concepts whose role in our conceptual scheme is pivotal in a way that makes them permanently resistant to definition and the drawing of stable boundaries.

2. Methodological and Other Preliminaries

When someone says causation is really "two kinds of thing," or a vague jumble of things, are they making a claim about our *concept* of causation, or a claim about the real-world

phenomena that causal talk is *directed* on? They might be making either claim, or both. The literature does not always distinguish between the options here.

Usually, the main focus is on "our concept" of causation. A pluralist might claim that, despite surface appearances, our talk of causation is guided on different occasions by different sets of criteria. In most cases, this will have the consequence that the worldly connections that count as causal will themselves be disunified (because they satisfy different criteria). But in principle, pluralism at the level of concepts allows the possibility that, at least in the actual world, all the cases that satisfy one set of criteria also satisfy the other. This correlation might even hold as a matter of natural law. In such a case, causation would be unified in the world but not in our thinking. (Compare Hesperus and Phosphorus, heat and molecular motion.)

Is the converse also possible? That is, perhaps we apply a univocal test when looking for causation, but the phenomena it picks out are, in some deep sense, plural in character. If so, causation would be unified in our thinking but not in the world.

(Compare the case of *jade*, at least according to the philosophers' story in which jade is two different minerals that we lump together into one category without realizing it.)

Views of causation that have this structure might seem natural to Cartesian dualists, who may say that physical and mental causation are deeply different phenomena that we get an informal handle on via a single abstract set of criteria. When this is labeled as a kind of "causal pluralism," it might be objected that if all the cases do satisfy one set of abstract criteria, then in one important sense causation is unified in the world as well.

In this survey I will try to be explicit about whether claims are being made about our words and concepts, on one side, or the phenomena that our thought and talk are aimed at, on the other. The total picture is something like this. We can distinguish three relevant sets of facts. One is the total set of facts about our habits of thought and talk in this area – how ordinary people, scientists, doctors, lawyers, and so on, use the word "cause" and its relatives, how responsibility is attributed and how explanations are given. A second set of facts concerns what the world contains and how the world runs. A third set of facts concerns the relations *between* the other two. For example, does the ordinary concept of cause succeed in picking out a real "natural kind"? Causal pluralism in its familiar forms usually makes claims about the first set of facts, those about causal

thought and talk. But the other sets of facts often become relevant, and it is important to keep an eye on exactly what sort of thing an alleged plurality is supposed to be.

A couple of other distinctions between pluralisms need to be made before moving on. The first is subtle. There are usually two steps in mainstream analyses of causation. Speaking metaphorically, one step is finding the right raw material, and the other is building the right structure out of it. We might decide that causation is essentially a matter of counterfactual dependence, but we then have to work out which kinds of counterfactual dependence suffice to make *C* a cause of *E*. Even bracketing the possibility of pluralism about the raw materials, philosophers may have often assumed too unified a view about how the raw materials should be put together (Hitchcock 2003). The usual target for philosophical analysis is what it is for one thing to *be a cause of* another. But there might be a family of causal concepts, including triggering, enabling, hastening, delaying, being linked in a causal chain.... Once we have worked out whether *C* hastened *E*, perhaps it is pointless to work out whether this is also enough for *C* to *be a cause of E*. Hastening is just what it is, and it is one genuine causal relation. The focus below will be primarily be on the question of pluralism at the level of raw materials, but sometimes the two questions may interact.

Lastly, there are some ways in which one might assert a plurality of causes that are not important here. Most views of causation will hold that any event will have multiple causes (for example, more immediate and more remote causes). This is compatible with each cause being a cause in the same sense (eg., a counterfactual sense), so it does not amount to pluralism of the kind under discussion. Secondly, an analysis of causation might make use of a *conjunction* of very different criteria. This usually counts as a univocal treatment of causation, unlike a disjunctive treatment, even though highly heterogeneous conjunctions should probably seem philosophically puzzling as well.

3. Two Concepts of Cause?

The simplest form of causal pluralism, and the one that has been developed in most detail, is the idea that there are two distinct *concepts* of cause that have become tangled together in our language. We have one term "cause" that tends to be guided by two different sets of criteria that pull apart in some cases. This kind of causal pluralism could

be fixed, in principle, by a re-labeling. If we began speaking of "cause₁" and "cause₂," no trouble would remain. Each of these concepts could be assessed in its own right, with respect to how it relates to the world. Each might pick out a natural kind, one might be empty, and so on. This low-key pluralism treats the word "cause" as analogous to ordinary examples of ambiguity, like "bank" or "bat." "Two concepts" proposals have been outlined in detail several times (Sober 1985, Hall 2004), and there have been a lot of informal suggestions along the same lines (always apparently with $n=2$).

The best worked-out view of this kind is seen in Hall (2004). Hall focuses on a particular family of tensions in many people's causal intuitions, which are often, and reasonably, taken to motivate a pluralist view of some sort or other (see also Hitchcock 1998, Schaffer 2000). I will refer to these as intuitions based on *difference-making* and *production* as rival marks of causation. Sometimes a factor with the right role in the physical production of an event will be classified as causal, whether or not it made a difference to the outcome. But sometimes what matters is whether the factor made a difference. These criteria can pull apart even in everyday cases. In cases of *redundant* causation we see the role of production intuitions. The bullet strikes its victim and produces his death. This is true regardless of whether other or not other bullets were right behind it, ready to produce the same effect. Being suitably located in a producing mechanism suffices for causation, apparently regardless of difference-making.

In cases of *causation by prevention* we see the role of difference-making intuitions. The air-traffic controller sees the two planes on collision course, but you prevent him from issuing a warning. Your action would usually be considered one cause of the ensuing crash. This is true despite the absence of a physical connection between your action and the effect. You made a difference to how things went, and difference-making suffices for causation in these cases.

Hall argues that the tension between production-based criteria and difference-making extends beyond intuitions about particular cases, to affect general principles. For example, many philosophers have wondered whether causation is transitive. Focusing on production makes it seem that the answer is yes; focusing on difference-making suggests that the answer is no. Hall's response to the situation is to say that we have two distinct concepts guided by different sets of criteria. There is causation in the sense of production,

and causation in the sense of difference-making. (Hall actually calls this one "dependence.") If we distinguish them, we find that many of the plausible general principles about causation apply to one kind of relation or the other, but not to both. For Hall and others, difference-making is naturally analyzed using counterfactuals, in the tradition of Lewis (1973), while production is harder to characterize.

In more recent work (unpublished), Hall has expressed doubts about the "two concepts" analysis. This is because there are cases where a factor C looks like a cause of E , despite passing neither of his two tests. These are cases where C is a *redundant threat-canceller*. C blocks a threat T that would otherwise have prevented E , but if C had not acted then D would have sufficed to block T instead. A threat-canceller like C can lack any physical connection to E (think of the air-traffic control case), and in this case C was not a difference-maker either. So C passes neither test but might reasonably be said to be a cause of E . (Here the reader may, with good reason, consider applying the ideas from Hitchcock sketched at the end of the previous section. Maybe there is no need to ask the standard "...was a cause of..." question in such a case.)

So the "two concepts" view has internal problems. And once we have moved into the pluralist camp, it is worth asking whether this approach is the most promising in general. If it were really the case that people had two concepts, then actual usage should contain certain kinds of evidence of this. We should not merely see attempts to *clarify* causal claims, but attempts to *switch* the hearer from one sense of "cause" to another. A good model would be provided by words like "mad" and "funny" which have two senses, related to each other but distinct, and which work in ways that sometimes make disambiguation requests appropriate. ("Did you mean it was funny as in *weird*, or as in *ha-ha*?") But this switching does not seem visible in causal discourse; there the situation seems more disorderly. That motivates a second kind of pluralism.

4. The Amiable Jumble

In 1984 Brian Skyrms, in a short paper about quantum mechanical puzzles, suggested that ordinary causal description is guided by an "amiable jumble" of criteria. These criteria usually work together in macroscopic cases, but can pull apart in special situations, especially those unearthed by physics. Causal thinking is a rough-and-ready

framework, well-suited to the everyday world, but not something that helps us describe the fundamental processes that make the world run.

The criteria Skyrms mentions as part of our "jumble" are familiar tools for analyzing causation, including spatio-temporal connection and instantiation of a regularity. Such a view will naturally add that there is no fact of the matter about which of these tests are the central ones; the evolution of the concept has not been of a kind that has made such a prioritization necessary. Different people are free to weight different tests differently, and free to use different weightings on different occasions. So the idea might be summarized as follows:

AMIABLE JUMBLE: " C was a cause of E " is true iff the relation between C and E satisfied some contextually appropriate combination of our amiably jumbled criteria for causation (instantiating a regularity, being spatio-temporally connected in certain ways, inducing counterfactual intuitions...).

Skyrms himself thought that the criteria tend to work together in everyday contexts, as noted above. His article was written prior to the focused work on counterfactual dependence that has revealed ways in which different criteria can pull apart in everyday cases (like cases of causation by prevention). It now seems that the jumble is not nearly as amiable as Skyrms envisaged; perhaps "cantankerous jumble" would be a more appropriate term.

This is a simple but sensible proposal, and related views have been defended by other writers. It is sometimes unclear which tools from the philosophy of language are most appropriate here, and how they are related. For example, how does the amiable jumble view relate to the idea that causation is a "vague" concept? Is this the same as claiming that causation is a "cluster concept"? What about the idea of "family resemblance"?

Many cases of vagueness have nothing to do with pluralism in the sense under discussion, because they involve the vague application of single test, often a one-dimensional gradient standard. (How many hairs can one have while still counting as bald?) But it seems reasonable to say that *one* way to get a kind of vagueness is for a

concept to be guided by a plurality of jumbled criteria that can pull in different directions. In the case of causation, many components of the cluster may themselves be vague.

Secondly, I take the amiable jumble proposal to be essentially the same as the claim that causation is a "cluster concept." Richard Healy, responding to the same quantum mechanical puzzles as Skyrms, proposed in 1994 that causation is a cluster concept, and the underlying picture is the same. These writers do not say much about how exactly a jumbled or cluster concept works. Longworth (forthcoming) remedies this with a detailed treatment of cluster concepts and their relation to causation. Modifying his account, we might say that a cluster concept is one whose application is guided by a set of distinct criteria where (i) none of the criteria are individually necessary, (ii) the entire set is clearly sufficient, (iii) at least one of the criteria must be met, and (iv) some proper subsets are sufficient, though many of these subsets generate marginal or contested applications.

The cluster-concept or amiable jumble view seems to do better than the two-concepts approach in accommodating everyday facts about the use of causal language. (Here I rely on informal observation, not on empirical data.) In particular, the cluster concept view does not predict that it should be possible to switch a hearer from one discrete "sense" of causation to another. What we might expect to see instead is a practice of guiding interlocutors to give more weight to some criteria and less weight to others, and an occasional willingness to give stipulative precisifications ("Well, what I mean by 'cause' here is..."). This is, I suggest, what we do tend to see.

If causation is a cluster concept, is this a case of "family resemblance"? In many ways this term seems fair, and Longworth (forthcoming) associates the two ideas closely. However, I reserve the term "family resemblance" for a more unusual option.

All the views discussed so far have drawn on *abstract* criteria that can be associated with causation. These include "instantiating a regularity," "being linked in a chain of counterfactual dependence," and so on. These criteria are not based on the particular physical character of a connection between two things. This might seem natural, as it seems unlikely that all the physical (and perhaps non-physical) connections that are recognized as causal by ordinary criteria will have similarities in their intrinsic make-up and structure. That fact usually pushes the philosopher towards a more abstract

treatment. But this push is not so forceful once a pluralist attitude is on the table. This brings us to the possibility of a more unorthodox form of analysis.

5. Anscombe, Minimalism, and Family Resemblance

Elizabeth Anscombe's "Inaugural Lecture" on causation (1971) raised, in a few striking paragraphs, the possibility of an analysis in which a range of familiar terms that we use to pick out various specific causal relations – terms like "scrape" and "burn" – are *semantically prior* to the general term "cause."

The word "cause" itself is highly general. How does someone show that he has the concept *cause*? We may wish to say: only by having such a word in his vocabulary. If so, then the manifest possession of the concept presupposes the mastery of much else in language. I mean: the word "cause" can be *added* to a language in which are already represented many causal concepts. A small selection: *scrape, push, wet, carry, eat, burn, knock over, keep off, squash, make* (e.g. noises, paper boats), *hurt*. But if we care to imagine languages in which no special causal concepts are represented, then no description of the use of a word in such languages will be able to present it as meaning *cause*. (1971 p. 93)

The last claim – that no word could mean *cause* in a language with no "special causal concepts" – is the key one. Anscombe, as I interpret her, is saying that the essential semantic role of the word "cause" in languages like English is to collect together a range of more specific relations that we also pick out with words like "squash," "push" and so on. The application that the word "cause" has to phenomena in the world goes *via* the application that these more specific terms have. If a language lacked any words like "scrape" and "push" there would be nothing in the language for the general term "cause" to collect together.

Anscombe did not give much more of an analysis than this, but one way to make the proposal more precise is to develop it as a kind of "minimalism" about causation.

MINIMALISM:

1. "*C* was a cause of *E*" is true iff the relation between *C* and *E* can also be described using some member of set *S*, or can be described as a chain of relations each of which can be described using some member of *S*.

2. S is a set of causal verbs and other linguistic formulas which represent "special causal concepts" in Anscombe's sense.

Initially, for the purpose of exploring the simplest version of such a view, we can imagine that the composition of set S is taken as unexplained. All we know is that some specific relations are grouped together by the word "cause."

The term "minimalism" is intended to suggest an analogy with minimalism about truth, which holds that the meaning of "true" is captured entirely by the T-schema: " p ' is true iff p " (Horwich 1990). For the minimalist, the word "true" is a device for adding and removing quotation marks, thereby making possible the compact expression of various claims that would otherwise require a convoluted or infinitely long sentence. The word "true" does not pick out some special matching or correspondence relation between signs and the world. Similarly, causal minimalism holds that to assert that some connection is causal is not to attribute some special empirical or modal feature to it. To call a connection causal is to say that this connection is one that can also be described using some unspecified member of S . We use the term "cause" to generalize and abstract, to form sentences like "Bill caused a lot of trouble," which have the same sort of role that "Everything Bill believes is true" has on a minimal theory of truth.

I do not know of an explicit defence of a minimalist view of this kind. Cartwright (1999, 2004), who draws on Anscombe, may be close to it, but her view may instead be closer to one of the other options discussed below.

Minimalism is philosophically intriguing, but has clear problems. First, it is significant that Anscombe's examples of "special causal concepts" describe relations between objects, not relations between events or facts. So the operation of abstraction that is achieved by using the word "cause" is not as simple as what we see in the case of minimalist theories of truth. It is hard to tell how damaging this is to the basic thrust of minimalism. A more important problem concerns the possibility of discovering novel kinds of causal relation. According to minimalism, no relation could be truly called causal if we had not previously developed a more specific description for that type of connection and grouped it as causal. So, for example, if minimalism was true it would make no sense for a person in 17th century England to wonder whether the Great Fire of

London had somehow caused the end of a plague epidemic that had preceded it, while having no idea of any mechanism by which this could have happened.

Can this problem be fixed while remaining within the spirit of Anscombe's view? It is at this point that the notion of family resemblance may be useful. Perhaps the relations grouped as causal are unified by a set of similarities between the relations themselves – between burning, pushing, scraping, and so on. These similarities might be real even if there is no abstract test that can be used to describe all of these relations (in terms of regularities or counterfactuals, for example). So when someone wonders whether some coincidence or sequence has an unknown causal relation underlying it, they are wondering whether there is *some* relation between the two events that, once understood, will be found to have a family resemblance to the familiar cases of causal relations.

Here I treat the family resemblance view as different from the amiable jumble or cluster concept view. The cluster concept view supposes that we apply a jumble of abstract tests; the family resemblance view appeals to similarities between the characteristics of the specific relations themselves.

To the extent that this is a distinct option, does it handle the phenomena of causal thought and talk as well as the amiable jumble view does? If we look at how people think about complicated causal relations, then the cluster view seems more promising. It seems clear that in many cases people do engage in an abstract assessment of difference-making, for example. This is so abstract a mode of assessment that it is hard to make sense of within an Anscombe-inspired approach. Recent empirical work suggests that abstract assessment of difference-making might be quite psychologically deep in us as well (Gopnik and Schulz 2004). On the other hand, there are simple observation-guided forms of everyday causal description which seem to be handled very naturally by Anscombe's view or something like it. I can *see* you squash the pie, and that is enough for me to say that you caused various things. This suggests that the right analysis might include a role for a cluster of abstract tests, and *and* a role for family resemblance at the level of the specific causal relations themselves.

To some this last move will also seem like a "kitchen sink" sort of option, philosophically. Is *anything* left out? Surely the pluralist urge can go too far. It can, but

perhaps what we should take from the preceding discussion is the idea that a cataloguing of tensions and complexities in our concept needs now to be accompanied by a different kind of analysis.

6. Causation as an "Essentially Contested" Concept?

The idea of an "essentially contested concept" (ECC) was introduced by W. B. Gallie (1956). Gallie's aim was to describe how some concepts remain permanently resistant to definition and analysis because of the pivotal place they occupy in our conceptual scheme. These are cases where it is not just hard to work out when the conditions of application of a term are met, but cases where the conditions for application themselves are, given the concept's role, permanently susceptible to being challenged and renegotiated.

I suggest that something like this analysis might be applied to the case of causation. The details of Gallie's analysis, which was designed to deal with cases like *art* and *democracy*, do not fit the case of causation very well. But his treatment provides a good starting point.

As Gallie originally conceived them, ECCs are used to appraise human activities and achievements. There are accepted paradigm cases that exert influence on how the term is used, but the complexity and changeability of the domain is such that there will be no obvious and undeniable rules for extending the application of the term to new cases. An ECC is used with evaluative loading, and there will typically be consequences when it is successfully applied. As a result, the term's normal use occurs against a background of recognized dispute over its application. Proper use of the term itself involves a contesting of other uses, and recognition that one's use will in turn be contested.

The concept of cause is not in this category. It is (mostly) used to describe relations between events or facts, not to appraise human activities. It also seems too strong to say that a recognition of the ongoing contesting of criteria is a background assumption of normal use. But the concept of cause has become embedded in our practices in such a way that this wrangling is very *likely*. And in some kinds of causal discussion, the concept can have the "aggressive" character that Gallie was interested in.

So "cause" (unlike "art" and "democracy") has low-level uses that are intended to be uncontroversial, but it also has a richer ECC-like role in some contexts.

We might then treat causation as a more low-key relative of an ECC – a non-accidentally contested concept? It is a concept that will be reliably subject to dispute with respect to its boundaries and criteria for application. We expect terms to acquire this role when their successful application has significant downstream consequences, but their domain is complex in ways that involve the absence of sharp borderlines that function as attractors to usage. As in the case of Gallie's ECCs, an accepted set of exemplars and a sense of a shared *purpose* behind diverse uses prevent a fragmentation into distinct concepts.

These ideas might be linked to tools developed in recent "inferentialist" philosophy of language (Brandom 2000). Terms like "freedom" and "cause" tend to have significant consequences when successfully applied, but complicated criteria for application. The result is frequent challenge to whatever "language entry rules" might be operating.

The concept of causation has these features as a consequence of its role in a family of important practices. Perhaps the crucial one is the assignment of responsibility. When it is established that a person is causally responsible for some event, they are often subject to praise, blame, and sanction. Within the standard set of problem cases in philosophy, cases of causation by omission are particularly relevant (Beebe 2004). Suppose you walk past a child who has fallen in a pond, who then drowns. We *might* say that you can be held responsible for the child's death even though (because of the absence of physical connection) you were not a cause of it. That is, we might insist that there can be a non-causal basis for moral responsibility. But there is a great deal of pressure in the other direction. It becomes more unproblematic to hold you responsible if we treat your act as a cause.

So Gallie's idea of an ECC, or a modification of it, might have useful application to the case of causation. This helps us move beyond the simple observation that the concept of causation seems to be disunified in complex ways, to an understanding of why this situation came about and how it is sustained.

Acknowledgment: I am grateful to Alison Gopnik, Thomas Hofweber, and Francis Longworth for helpful discussions of these issues.

References

- Beebe, H. (2004). "Causing and Nothingness." In Collins, Hall, and Paul (2004), pp. 291-308.
- Anscombe, G. E. M. (1971). "Causality and Determination," reprinted in E. Sosa and M. Tooley (eds.), *Causation*. Oxford: Oxford University Press, 1993, pp. 88-104.
- Brandom, R. (2000). *Articulating Reasons: An Introduction to Inferentialism*. Cambridge MA: Harvard University Press.
- Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.
- Cartwright, N. (2004). "Causation: One Word Many Things." *Philosophy of Science* 71: 805–819.
- Gallie, W. B. (1956). "Essentially Contested Concepts." *Proceedings of the Aristotelian Society* 56: 167-98.
- Collins, J., E. Hall, and L. Paul, (eds.) (2004). *Causation and Counterfactuals*. Cambridge MA: MIT Press.
- Gopnik, A. and L. Schulz. (2004). "Mechanisms of Theory-Formation in Young Children." *Trends in Cognitive Sciences*. 8: 371-377.
- Hall, E. (2004). "Two Concepts of Causation," in J. Collins, E. Hall, and L. Paul, (eds.) *Causation and Counterfactuals*. Cambridge MA: MIT Press, pp. 225-276.
- Healey, R. (1994). "Nonseparable Processes and Causal Explanation." *Studies in History and Philosophy of Science* 25: 337-374.
- Hitchcock, C. (1998). "Causal Knowledge: That Great Guide of Human Life." *Communication and Cognition* 31: 271-296.
- Hitchcock, C. (2003). "Of Humean Bondage." *British Journal for the Philosophy of Science* 54: 1–25.
- Horwich, C. (1990). *Truth*. Oxford: Clarendon Press.
- Lewis, D. K. (1973). "Causation." *Journal of Philosophy* 70: 556-67.

Lewontin, R. (1974). "The Analysis of Variance and the Analysis of Causes." *American Journal of Human Genetics* 26: 400-411.

Longworth, F. (forthcoming). "Is Causation a Cluster Concept?"

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

Skyrms, B. (1984) "EPR: Lessons for Metaphysics," in P. French, T. Uehling, and H. Wettstein *Midwest Studies in Philosophy IX*. Minneapolis: University of Minnesota Press, pp. 245-255.

Sober, E. (1985). "Two Concepts of Cause." In P. Asquith and P. Kitcher (eds.) *PSA 1984, Volume II*. East Lansing: Philosophy of Science Association pp. 405-24.