




2010

## Mechanisms (Oxford)

Stuart Glennan

*Butler University*, [sglennan@butler.edu](mailto:sglennan@butler.edu)

Follow this and additional works at: [http://digitalcommons.butler.edu/facsch\\_papers](http://digitalcommons.butler.edu/facsch_papers)

 Part of the [Epistemology Commons](#), [Philosophy of Mind Commons](#), [Philosophy of Science Commons](#), and the [Social and Behavioral Sciences Commons](#)

---

### Recommended Citation

Mechanisms (2010). *The Oxford Handbook of Causation*, edited by Helen Beebee, Christopher Hitchcock and Peter Menzies. Oxford: Oxford University Press

This Article is brought to you for free and open access by the College of Liberal Arts & Sciences at Digital Commons @ Butler University. It has been accepted for inclusion in Scholarship and Professional Work - LAS by an authorized administrator of Digital Commons @ Butler University. For more information, please contact [omacisaa@butler.edu](mailto:omacisaa@butler.edu).

**Mechanisms**  
Stuart Glennan  
Butler University

Forthcoming in 2009 in the Oxford Handbook of Causation

## **1. Introduction**

The term ‘mechanism’ is commonly used to refer to a variety of systems or processes that produce phenomena in virtue of the arrangement and interaction of a number of parts. The term was originally applied to products of human design like watches or water wheels, but has, at least since the seventeenth century been used equally to describe systems (like cells) or processes (like those that produce sunspots) that are of natural origin (Dijksterhuis 1969).

Mechanism is undoubtedly a causal concept, in the sense that ordinary definitions and philosophical analyses explicate the concept in terms of other causal concepts such as production and interaction. Given this fact, many philosophers have supposed that analyses of the concept of mechanism, while they might appeal to philosophical theories about the nature of causation, could do little to inform such theories. On the other hand, methods of causal inference and explanation appeal to mechanisms. Discovering a mechanism is the gold standard for establishing and explaining causal connections. This fact suggests that it might be possible to provide an analysis of causation that appeals to mechanisms.

There have been a variety of attempts to explicate the concept of mechanism and to deploy it in understanding causation and causal explanation. Salmon (1984) and Dowe (2000) treat mechanisms as a nexus of continuous physical processes. Bechtel and Richardson (1993), Glennan (1996) and Machamer, Darden and Craver (2000) treat

mechanisms as systems of interacting parts. As the process approach is discussed elsewhere (Chapter 10), I will focus on the mechanical systems approach, which for brevity I shall refer to simply as the mechanical approach.

Glennan, who is the most explicit in trying to develop a mechanical account of causation, characterizes mechanisms in this way:

A mechanism underlying a behavior is a complex system which produces that behavior by the interaction of a number of parts according to direct causal laws (Glennan 1996: 52).

Glennan then suggests that two events are causally related when and only when they are connected by an intervening mechanism.

I will explore three issues surrounding the adequacy of the mechanical approach to causation. First, I consider whether the appeal to laws or invariant generalizations in characterizing interactions between parts of mechanisms either makes the mechanical theory circular or reduces it to a regularity, counterfactual or manipulability theory. Second, I discuss Machamer, Darden and Craver's argument that the proper understanding of the causal productivity of mechanisms requires the recognition of the novel metaphysical category of *activities*. Third, I discuss the relationship between mechanical theories and process theories.

## **2. Laws, Generalizations and Mechanisms:**

Machamer, Darden and Craver (hereafter MDC) and Woodward (2002) have raised concerns about Glennan's use of laws to characterize interactions between parts of mechanisms. According to one common understanding of scientific laws, laws must be exceptionless regularities of unrestricted scope. If one understands laws in this way, it is

implausible to assume that all interactions between parts of mechanisms are law-governed, because the regularities involved in the operation of mechanisms can be fragile and subject to exceptions and breakdowns.

But this objection is largely an issue of terminology. Glennan explicitly adopts an alternative usage which understands laws to comprise a large class non-accidental, counterfactual supporting generalizations, many of them of restricted scope.<sup>1</sup> For example, according to this usage Mendel's law of independent assortment really is a law, even though it is subject to exceptions (as when loci lie near to each other on the same chromosome). While this usage has the virtue of stressing that the honorific 'law' is applied to many generalizations of this character, Glennan has more recently (2002) borrowed terminology from Woodward (2000), describing such generalizations as direct invariant change-relating generalizations. These generalizations characterize interactions in which changes in the properties of one or more parts bring about changes in a property of another part. The requirement that these generalizations be direct dictates that the change of properties in one part brings about the change in the properties of another part without changing the properties of another part that is intermediate within a causal chain.<sup>2</sup>

Whatever one calls generalizations describing interactions between parts of mechanisms, critics can argue that that these generalizations do all the causal work. A mechanist must have some account of truth conditions for these generalizations, and the suspicion is that any real insight into the nature of causation will lie in these conditions rather than in the analysis of mechanisms. In particular, given the appeals that Glennan (2002) and Craver (forthcoming) make to Woodward-style counterfactuals, one might suspect that the mechanical account will reduce to a manipulability account.

The mechanist can respond to this objection by appealing to the hierarchical character of mechanisms. Mechanisms consist of parts whose direct interactions are productive of some behavior, but what counts as a direct interaction depends upon the level of analysis. In general, the parts of mechanisms are themselves mechanisms whose behaviors depend upon the organization and interaction of the parts of those parts. So for instance, one can explain the causal processes controlled by neuron circuits by invoking direct interactions between neurons at synapses. But the mechanism of synaptic transmission is itself a complex one, involving many constituent parts. These parts (e.g., vesicles, neurotransmitters, ions and ion channels) are themselves mechanisms composed of parts whose interactions explain their behavior. Thus, one can see that the change-relating generalizations--in this case a generalization describing how one action potential stimulates or inhibits another action potential--are explained by appeal to underlying mechanisms. These generalizations are *mechanically explicable* (Glennan 1996: 61-63). Thus, the fact that such generalizations are appealed to in characterizing a causal relation does not undermine the mechanical character of the account.

The manipulability and mechanist accounts provide different criteria for identifying causal connections. On the one hand, there is the manipulability criterion--in this case, if one could manipulate the pre-synaptic neuron one would stimulate or inhibit the post-synaptic neuron. On the other hand, there is the criterion of an identifiable mechanism--in this case one has identified the mechanism of synaptic transmission. Is there a reason to think that one criterion more genuinely captures the nature of the causal relation?

Proponents of the manipulability theory will immediately point to what appears to be the Achilles heel of the mechanical account. Unless there is an infinite downward chain of nested mechanisms, sooner or later one will run out of mechanisms. One reaches a point where the parts of mechanisms interact in accordance with change relating generalizations that are not mechanically explicable. To take a simple case, imagine a Newtonian world in which two particles accelerate towards each other in virtue of their gravitational attraction. This attraction is characterized by a change-relating generalization describing the particles' changes in velocity, but there is not (we'll suppose) a mechanism that explains this. It is just a brute fact that massive objects *cause* other objects to accelerate toward them. Furthermore, the critic can make the following sort of reductive argument: Because analysis of mechanically explicable generalizations at higher levels will ultimately bottom out in mechanisms whose parts interact in non-mechanically-explicable ways, it will ultimately be the case that the causal character of such generalizations depends upon on what can be called fundamental or mechanically inexplicable generalizations. The virtue of the manipulability account is that it seems to work even at the fundamental level. While, at the fundamental level I cannot find a mechanism connecting *X* and *Y*, I can at least determine whether there is a causal connection by seeing if I can wiggle *X* to wiggle *Y*. In light of this objection, the manipulability theorist might argue that mechanists need manipulability but manipulability theorists can live without mechanisms.

But the advantage of the manipulability theory is less than one supposes. At some level the criticism of the mechanistic theory is that it cannot reductively eliminate causal concepts. But, as Woodward (2003) takes pains to point out, neither can the

manipulability theory. At the root of the manipulability analysis is the transparently causal notion of an intervention. At the level of fundamental causal connections, how does one know that one has wiggled  $X$  (and just  $X$ )? One can have some confidence that one is wiggling  $X$  if one has an account of the wiggling mechanism, as well as an account of what mechanisms are used to isolate  $X$  from other factors that might be causally connected to  $Y$ . In the absence of such an account, the manipulability theorist has only the observation that the wiggling of  $X$  was followed by the wiggling of  $Y$  and the *supposition* that the wiggle of  $Y$  depended on that of  $X$ .

Minimally, the manipulability theorist has made an important point about the epistemology of causation. It is frequently the case that one can establish with reasonable certainty that a variable is causally relevant to another without knowing anything about the mechanism by which the variables are connected. Experimental manipulations can provide evidence *that* variables are connected, even in the absence of mechanical knowledge of *how* they are connected. But this epistemologically important point does not legitimate the manipulability theory as a metaphysical account of causation. To see this, consider the truth conditions for the interventionist counterfactuals that underlie the manipulability theory. Do these truth conditions depend on singular or general facts? Woodward's answer is clear. Although causal claims are often made at the type level, the truth of these claims depends upon the truth of claims about individuals:

a claim such as ' $X$  is causally relevant to  $Y$ ' is a claim to the effect that changing the value of  $X$  instantiated in particular individuals will change the value of  $Y$  located in particular individuals (Woodward 2003: 40).

So, to understand the truth conditions for type-level claims, we must understand the truth conditions for the token claims. These truth conditions are specified in terms of interventionist counterfactuals. If one were to intervene on (this instance of  $X$  to change its value) one would change  $Y$ . But how are we to understand the truth conditions for these counterfactuals? Suppose at some time  $t$  I flip a switch ( $e_1$ ) and a light goes on ( $e_2$ ). As an epistemic matter, I establish that  $e_1$  caused  $e_2$  by manipulating the switch at other times  $t'$ , but the truth of the causal claim at  $t$  does not depend upon the truth of correlations at  $t'$ . It depends upon the singular counterfactual dependency at  $t$ .

Stathis Psillos (2004) has raised this concern about Woodward's account, arguing that Woodward has given evidence conditions for interventionist counterfactuals, but that he hasn't given truth conditions. Psillos suggests that the appropriate way to correct this deficiency is to appeal to laws of nature. Although he does not spell out how this appeal would work in detail, roughly it must be that the counterfactual dependency would be true in virtue of the fact that the sequence of antecedent and consequent was a tokening of a lawful generalization. To make this proposal work, one needs an account of laws of nature, and Psillos endorses the Mill-Ramsey-Lewis (MRL) approach, in which laws are taken to be generalizations that are theorems (or axioms) of a deductive system that provides the best combination of simplicity and strength (Lewis 1973: 73).

In adopting this view, Psillos has just traded one problem for another. The MRL account of laws is a Humean position that grounds the truth of laws in the totality of particular facts in the actual world (cf. Chapter 7). If truth conditions for counterfactuals are grounded in MRL laws, Psillos' proposal does provide truth conditions for counterfactuals that are grounded in epistemically accessible facts about the actual world,



but then the facts upon which a singular counterfactual (like the fact that if I were to flip the switch I would turn on the light) depends are facts not just about the particular instance but about the whole class of particulars. Thus, it violates the principle that the truth conditions for singular causal claims should be intrinsic, that whether a change of an instance of *X* causes a change of an instance of *Y* should depend only upon the local facts surrounding these instances. Thus Woodward should rightly reject Psillos' attempt to ground the truth of singularist interventionist counterfactuals in MRL laws.

Except in the case of fundamental laws, the mechanical theory of causation can provide the truth conditions for interventionist counterfactuals that the manipulability theory seems to lack. When causal relations are mechanically explicable, what makes it the case that wiggling some *X* will produce the wiggling of some *Y* is that there is an intervening mechanism between *X* and *Y*.<sup>3</sup> Understanding the nature, structure and functional organization of the parts that make up that mechanism will allow one to determine the range of counterfactual circumstances under which the dependency between *X* and *Y* would be maintained--roughly those circumstances in which the mechanism will not break down.

Given that any variables connected by intervening mechanisms depend on mechanically inexplicable causal connections, it does seem incumbent on the mechanical theorist to say something how one can know that such connections obtain. The mechanist supposes that causal connections are constituted by interactions between parts of mechanisms, and that the interactions between parts of mechanisms may, on further analysis, be constituted by nested mechanisms, but that ultimately one will bottom out with parts that interact, but where these interactions aren't mechanically explicable.

Suppose that one observes that changes in one part are followed by changes in another part--perhaps changes that occur without intervention or perhaps changes that one believes to be caused by an intervention. How does one know that the changes in one part really *produce* the changes in the other part? This is Hume's problem reproduced at the level of the most basic constituents of mechanisms. Following the MRL approach, one might take the fundamental laws to be those change-relating generalizations that give us the simplest and strongest account of the facts in the world. Neither the mechanist nor the manipulability theorist needs to suggest that fitting within the strongest and simplest account is what makes it true that these variables are causally connected, but the mechanist can say that this is the best reason we have, at the fundamental level, to accept certain generalizations as expressing causal relations.

### **3. Activities and Interactions**

MDC's (2000) analysis of mechanisms is distinguished by its introduction of the concept of an activity. MDC characterize mechanisms as 'entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions' (Machamer, Darden, and Craver 2000: 3). MDC claim that the introduction of the term 'activity' reflects a crucial metaphysical innovation. They argue for this claim by characterizing the debate over the nature of mechanisms in terms of a dispute between three ontological stances. First there are the substantialists (including Bechtel and Richardson and Glennan) whose basic ontology is an ontology of entities. 'Substantialists confine their attention to entities and properties, believing it is possible to reduce talk of activities to talk of properties and their transitions' (MDC 2000: 4). Second, there are process ontologists (notably Rescher (1996)), who 'reify activities and

attempt to reduce entities to processes' (*ibid*). They describe their new third way as 'dualist', requiring both entities and activities as ontologically irreducible categories.

MDC appear to overstate the novelty of their position because they fail to recognize the dualist character of the substantialist account. Where MDC speak of entities and activities, Bechtel, Richardson and Glennan speak of parts and interactions. No account of a mechanism can get along talking about the entities alone, but for Glennan and for Bechtel and Richardson this additional ingredient is provided by interactions.

MDC argue that this language is insufficient:

[I]t is artificial and impoverished to describe mechanisms solely in terms of entities, properties, interactions, inputs-outputs, and state changes over time. Mechanisms do things. They are active and so ought to be described in terms of the activities of their entities, not merely in terms of changes in their properties (MDC, 5).

MDC equate interactions with state changes over time. What this suggests is that the real target of their criticism is a Humean conception of interactions in which interactions are characterized 'merely in terms of changes in ... properties.' Their view is that productivity cannot be reduced to mere change in properties. Given the manifold problems with Humean approaches to causation, this is certainly a respectable position to take, but neither Glennan nor Bechtel and Richardson adopt the Humean approach. If activities can be productive, so can interactions.

Machamer (2004) and Bogen (2005) have recently offered more detailed defenses of the metaphysical significance of activities. The crux of their defense is that activities

provide a better and more natural approach to understanding productivity than do interactions, because activities can help us understand productivity without appeal to counterfactuals.

First consider how counterfactuals appear to give some understanding of the productive character of interactions. When a change in one part of a mechanism produces a change in another part of a mechanism, the claim that the first change *produces* rather than just precedes the second change is (ignoring a lot of details) taken to be licensed by whatever evidence one has for the counterfactual claim that if the first change had not occurred, neither would have the second. The modal character of the productivity is cashed out in terms of the counterfactual.

Bogen worries however that this analysis makes the truth of a causal claim depend upon what would have happened in other circumstances--what he calls counterfactual generality--rather than upon what actually does happen. At root, both Bogen and Machamer are appealing to a well-known argument from Elizabeth Anscombe:

If A comes from B, this does not imply that every A-like thing comes from some B-like thing or set-up or that every B like thing or set-up has an A-like thing coming from it; or that given B, A had to come from it, or that given A, there had to be B for it to come from. Any of these may be true, but if any is, that will be an additional fact, not comprised in A's coming from B (Anscombe 1993: 92).

. Anscombe insists that the production of a particular B from a particular A is independent of any actual or counterfactual regularity. Bogen and Machamer seem to

think that activities capture the notion of productivity without reference to any such regularities.

I will not dispute the claim that the productive continuity between cause and effect is a fact about the actual world in the particular case. What I will dispute is that Anscombe, Machamer or Bogen have adequately explained what this productivity is. Anscombe writes ‘causality consists of the derivativeness of an effect from its causes. This is the core, the common feature, of causality in its various kinds. Effects derive from, arise out of, come of, their causes’ (*ibid*). MDC echo this when they characterize activities: ‘Activities are the producers of change. They are constitutive of the transformations that yield new states of affairs or new products’ (MDC 2000: 4). But these characterizations of activities and causes do little more than offer a set of synonyms for ‘cause’. They hardly elucidate the concept of cause or production. Neither Anscombe nor MDC would worry too much about this, because they believe that we acquire our understanding of causation by starting with the understanding of particular activities--pushing, bending, hitting, etc. But while this is certainly true, to the extent that our concept of cause does mean something, we must give some account of what it means in general, and to provide a synonym for ‘cause’, like ‘produce’, is not to provide an analysis. The virtue of the manipulability approach to causation is that it tells one something quite general about causes--that they can be used, in an idealized way--to manipulate effects. The virtue of the mechanical approach is again that it tells one something quite general about causes--that causes and effects will generally be connected by intervening mechanisms. Neither of these theories provides reductive analyses of causation, but both say something non-trivial about the nature of causation.

## 4. Mechanical Systems and Mechanical Processes

While the term ‘mechanism’ is currently most commonly associated with the concept of mechanism identified with the work of Bechtel, Glennan, and Machamer, Darden and Craver, the term ‘mechanism’ first entered into the contemporary discussion of causation and explanation with the work of Salmon and Railton, and the ‘causal-mechanical’ approach to explanation (Railton 1978; Salmon 1984). What distinguishes the earlier Salmon/Railton approach from the more recent approach is that Salmon and Railton conceive of mechanisms as a network of interacting *processes*, whereas the more recent mechanists think of mechanisms as *systems* — organized collections of parts.

To help understand the difference between the approaches, consider paradigm cases of each sort of mechanism. For a systems theorist a paradigm case might be a toilet, while for the process theorist a paradigm case might be a baseball striking a window. The toilet is a *thing*—a structured system consisting of parts (valves, levers, floats, etc.) that interact in regular ways. A baseball striking a window involves a process which involves a series of events (the pitch of the baseball, the hit, the collision with the window, etc.), but we can’t think of this sequence of events as a thing. The *operation* of system mechanisms give rise to processes (e.g., toilets flush), but these processes are, in virtue of the stability of the mechanism, regular and repeatable.

By attending to the distinction between these two kinds of mechanisms, we can see a significant limitation in the mechanical approach to causation proposed in Glennan 1996. The hope expressed in that paper was that any two causally connected events that were not connected in virtue of a fundamental (i.e., mechanically inexplicable) law, would be connected by the operation of an intervening mechanism. If mechanisms are construed as systems as they are in the accounts of Glennan, Bechtel and MDC, it is

evident that there are many true singular causal claims in which a cause is not connected to its effect via the operation of a mechanism qua system. When the baseball breaks the window, there is a causal process that leads from the pitch of the baseball to the breaking of the window, but there is no mechanism qua system for window breaking.

Clearly one of the attractions of the process approach is that it appears to give an account of causal connectedness for singular causal sequences. But, it does so at a cost. As Hitchcock (1995) has argued, process theories appear unable to provide any intuitive understanding of the explanatory relevance of causes for effects. Relatedly, a process theory that identifies interactions in terms of exchange of conserved quantities does not provide informative and explanatory accounts of the nature of interactions between parts at higher levels of organization (Glennan 2002; Psillos 2004).

What appears to be required to provide a more adequate and general mechanistic account of singular causal sequences is to describe causal processes in terms of the interactions of parts, as is done in the systems approach, but to recognize that in many singular causal chains, the parts are not organized in a stable system. Thus for instance, one can conceive of the baseball sequence in terms of the interaction of a set of parts (baseball, bat and window), but these parts are not arranged in any stable configuration. Consequently, the next pitch won't necessarily be hit, and if hit, it won't likely travel on exactly the same trajectory. But what one does know is that when bats hit balls in certain ways (i.e., with certain velocities and spins) they alter trajectories in reliable ways that can be described by 'change-relating generalizations'. Thus we see the striking of a ball in a certain way as explanatorily relevant for the changing of that ball's velocity and ultimately for the breaking of the window. The singularity of the causal sequence arises

from the fact that the particular arrangement and state of the parts of the mechanism on a given occasion are ephemeral, but the parts themselves and their behavioral dispositions are robust. An account of such ‘ephemeral mechanisms’ is briefly suggested in Glennan 2002, but much work would be needed to show how such an approach could improve upon existing accounts of singular causation.

## **5. Further Reading**

For a general introduction to recent work on mechanism it is good to begin with Glennan 1996 and MDC 2000. The former work is more focused on issues of causation, while the latter discusses a range of topics, including models, explanation and discovery. Bechtel and Richardson 1993 is a book-length treatment of approaches to discovering mechanisms, especially in biological systems, but it anticipates much of the general mechanistic approach developed over the last decade. It is helpful to contrast this view of mechanism with the theory of causal processes discussed in chapter 10 of this volume.

Concerning activities, the reader might begin with Machamer’s (2004) defense. Tabery 2004 argues that both MDCs activities and Glennan’s interactions are necessary for a proper account of mechanisms. Psillos 2004 provides a useful analysis and critique of activities as well as of Glennan’s approach, and seeks to provide an account of the relationship between mechanistic and manipulationist and counterfactual accounts of causation.

There has been considerable work done in recent years on developing a theory of causal-mechanical explanation. In the 1980s and 1990s, the term ‘causal mechanical explanation’ was chiefly applied to Salmon’s (1984) approach. Glennan 2002 contrasts Salmon’s approach to one consistent with the systems approach to mechanism. Thagard



1999 provides a related account of mechanisms and mechanistic explanation in the context of medical science. Bechtel and Abrahamsen 2005 offers a general account of mechanistic explanation, contrasting it with nomological approaches. Using case studies from neuroscience, Craver forthcoming provides a comprehensive theory of mechanistic explanation. Craver's book also has substantial discussions of the relationship between mechanisms, manipulability and causation and of the implications of the mechanistic approach for reduction and problems of higher-level causation.

## **Notes**

<sup>1</sup> Advocates of this usage of the term 'law' include Cartwright (1999) and Mitchell(1997).

<sup>2</sup> The notion involved here is essentially the same as Woodward's notion of a direct cause (2003), in which direct causes are those causes which can bring about effects while holding all other factors fixed. Directness is relativized to a particular decomposition of a system into parts.

<sup>3</sup> This view is not really at odds with Woodward's, because the path diagrams Woodward uses to represent causal structure are naturally interpreted as diagrams of mechanisms.

## References

- Anscombe, G. E. M. (1993). 'Causality and Determination', in Ernest Sosa and Michael Tooley(eds.), *Causation*. Oxford: Oxford University Press, 88-104.
- Bechtel, William, and Adele Abrahamsen (2005). 'Explanation: A Mechanist Alternative', *Studies in the History and Philosophy of Biology and the Biomedical Sciences* 36(2): 421-441.
- Bechtel, William, and Robert C. Richardson (1993). *Discovering Complexity : Decomposition and Localization as Strategies in Scientific Research*. Princeton, NJ: Princeton University Press.
- Bogen, Jim (2005). 'Regularities and Causality: Generalizations and Causal Explanations', *Studies in the History and Philosophy of Biology and Biomedical Sciences* 36(2): 3970420.
- Cartwright, Nancy (1999). *The Dappled World : A Study of the Boundaries of Science*. Cambridge, U.K. ; New York: Cambridge University Press.
- Craver, Carl F. (forthcoming). *Explaining the Brain: What a Science of the Mind-Brain could be*. New York: Oxford.
- Dijksterhuis, E. J. (1969). *The Mechanization of the World Picture*. A galaxy book GB298. [Mechanisering van het wereldbeeld.]. London, New York: Oxford University Press.
- Dowe, Phil (2000). *Physical Causation*. Cambridge studies in probability, induction, and decision theory. Cambridge ; New York: Cambridge University Press.
- Glennan, Stuart S. (1996). 'Mechanisms and the Nature of Causation', *Erkenntnis* 44(1): 49-71.

- . (2002). 'Rethinking Mechanistic Explanation', *Philosophy of Science* 69(3 Supplement): S342-S353.
- Hitchcock, Christopher R. (1995). 'Discussion: Salmon on Explanatory Relevance', *Philosophy of Science* 62(2): 304-320.
- Lewis, David K. (1973). *Counterfactuals*. Oxford: Blackwell.
- Machamer, Peter (2004). 'Activities and Causation: The Metaphysics and Epistemology of Mechanisms', *International Studies in the Philosophy of Science* 18(1): 27-39.
- Machamer, Peter, Lindley Darden, and Carl F. Craver (2000). 'Thinking about Mechanisms', *Philosophy of Science* 67(1): 1-25.
- Mitchell, Sandra D. (1997). 'Pragmatic Laws', *Philosophy of Science* 64(4): S468-S479.
- Psillos, Stathis (2004). 'A Glimpse of the Secret Connexion: Harmonizing Mechanisms with Counterfactuals', *Perspectives on Science* 12(3): 288-319.
- Railton, Peter (1978). 'A Deductive-Nomological Model of Probabilistic Explanation', *Philosophy of Science* 45: 206-226.
- Rescher, Nicholas (1996). *Process Metaphysics : An Introduction to Process Philosophy*. SUNY series in philosophy. Albany: State University of New York Press.
- Salmon, Wesley C. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton, N.J.: Princeton University Press.
- Thagard, Paul (1999). *How Scientists Explain Disease*. Princeton, N.J.: Princeton University Press.
- Woodward, James (2000). 'Explanation and Invariance in the Special Sciences', *British Journal for the Philosophy of Science* 51(2): 197-254.

- . (2002). 'What is a Mechanism? A Counterfactual Account', *Philosophy of Science* 69(3 Supplement): S366-S377.
- . (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.