

# Rethinking Mechanistic Explanation

Stuart Glennan<sup>†‡</sup>

Butler University

---

Philosophers of science typically associate the causal-mechanical view of scientific explanation with the work of Railton and Salmon. In this paper I shall argue that the defects of this view arise from an inadequate analysis of the concept of mechanism. I contrast Salmon's account of mechanisms in terms of the causal nexus with my own account of mechanisms, in which mechanisms are viewed as complex systems. After describing these two concepts of mechanism, I show how the complex-systems approach avoids certain objections to Salmon's account of causal-mechanical explanation. I conclude by discussing how mechanistic explanations can provide understanding by unification.

---

**1. Introduction.** Recent philosophical work on scientific explanation has been dominated by two approaches: Salmon's (1984) causal-mechanical approach suggests that scientific explanations explain events by showing how they fit into the causal structure of the world, while Kitcher's (1989) explanatory unification approach suggests that explanations explain by showing how diverse phenomena fit into unifying patterns. While both Salmon and Kitcher acknowledge that many explanations unify and that many other explanations exhibit causal mechanisms, they differ sharply on which of these approaches is fundamental. Salmon claims that unification occurs because many physical processes utilize the same basic causal mechanisms, while Kitcher suggests that our claims about causes are ultimately based upon unifying explanatory patterns. As Kitcher puts it, "the unification view . . . proposes to ground causal claims in claims about explanatory dependency rather than *visa versa*" (Kitcher 1989, 436).

<sup>†</sup>Send requests for reprints to the author, Department of Philosophy and Religion, 4600 Sunset Ave., Butler University, Indianapolis, IN 46208-3487; e-mail: [sglennan@butler.edu](mailto:sglennan@butler.edu).

<sup>‡</sup>Thanks to Lindley Darden for helpful comments on this paper. A longer version of this paper with considerably fuller references is available on-line at <http://scistud.umkc.edu/psa2k/>.

Philosophy of Science, 69 (September 2002) pp. S342-S353. 0031-8248/2002/69supp-0029\$10.00  
Copyright 2002 by the Philosophy of Science Association. All rights reserved.

While I believe that Salmon is ultimately correct, he has been hampered in defending his view by an inadequate analysis of the nature of causal mechanisms. My aim in this paper will be to provide an improved analysis of mechanisms—one adequate to illuminate the unifying power of mechanistic explanation.

**2. Salmon and Railton on Causal-Mechanical Explanation.** Peter Railton (1978) first introduced the idea of mechanism into the contemporary literature on explanation with his deductive-nomothetic model of probabilistic explanation (DNP). The DNP model was meant to be an alternative to Hempel's inductive statistical (IS) model. Railton was concerned with Hempel's requirement that the explanans of an IS explanation render the explanandum probable or nomically expectable. On Railton's view, explanations describe causes, and sometimes the sequence of events leading up to the event to be explained may be improbable. According to Railton, while an explanation of some event may include a reference to a law that renders the event nomically expectable, the account must be supplemented by "an account of the mechanism(s) at work" (1978, 748). Railton is deliberately vague on just what a mechanism is, indicating only that an "account of the mechanism(s)" is "a more or less complete filling-in of the links in the causal chains" (1978, 748).

Salmon's work on causal-mechanical explanation can be seen as an attempt to elaborate on Railton's earlier account of mechanistic explanation. Curiously, nowhere in this book does Salmon offer an explicit definition of mechanisms. Rather, he provides a characterization of what he calls the "causal nexus," which he takes to be a vast network of interacting causal processes. Much of Salmon's work is directed at providing an adequate definition of just what count as causal processes and interactions. According to his definition, a process is an entity that maintains a persistent structure through space-time, a causal process is a process capable of transmitting changes in its structure, and a causal interaction is an intersection between causal processes in which an alteration of the persistent properties of those processes occurs. In Salmon's original formulation interactions were defined in terms of a counterfactual criterion of mark transmission, but in more recent versions (1994) he has eliminated reference to counterfactuals, defining causal interactions in terms of exchange of conserved quantities. While Salmon has spent considerable effort subsequent to 1984 trying to refine his characterizations of processes and interactions, his view of the causal nexus as a network of interacting processes remains essentially unchanged. Causal-mechanical explanation exemplifies what Salmon calls the ontic conception of explanation. Explanations are not arguments, but are rather descriptions of features of a mind-independent reality—the causal structure of the world.

**3. The Complex-Systems Approach to Mechanism.** While the term “causal-mechanical explanation” has come in the literature to refer to the theories of explanation proposed by Railton and Salmon, a number of philosophers have been developing a quite different analysis of mechanism—one that leads to an alternative conception of causal-mechanical explanation. Central to this alternative conception is the idea that a mechanism is a “complex system.” The germ of this conception can be found in Wimsatt’s work, which in turn has antecedents in the work of Herbert Simon. This conception has been developed to a greater degree in my own work (Glennan 1996), as well as that of Bechtel and Richardson (1993) and, most recently, Machamer, Darden, and Craver (2000). While there are differences between these approaches, they all share the view that mechanisms are complex systems.

A number of definitions of mechanism have been suggested, but my preferred one is as follows:

- (M) A mechanism for a behavior is a complex system that produces that behavior by the interaction of a number of parts, where the interactions between parts can be characterized by direct, invariant, change-relating generalizations.

I shall not repeat arguments I have made elsewhere (Glennan 1996) for this definition, but a few clarifications are in order. First, mechanisms are not mechanisms *simpliciter*, but mechanisms *for* behaviors. A complex system may exhibit several different behaviors, and the decomposition of the system will depend upon which behavior is under consideration. For instance, a heart is both a mechanism that pumps blood and a mechanism that produces noise; the heart qua pump may admit of a different decomposition than the heart qua noisemaker. Second, mechanisms consist of a number of parts. These parts must be objects, in the most general sense. They must have a relatively high degree of robustness or stability; that is, in the absence of interventions, their properties must remain relatively stable. Generally these parts can be spatially localized.

A mechanism operates by the interaction of parts. An interaction is an occasion on which a change in a property of one part brings about a change in a property of another part. For instance, a change in the position of one gear within a clock mechanism may bring about the change in the position of an interlocking gear. “Interaction” is a causal notion that must be understood in terms of the truth of certain counterfactuals. The stipulation that these interactions can be characterized by invariant, change-relating generalizations is meant to capture the relevant counterfactual truth claims.

I borrow the phrase “invariant, change-relating generalization” from Jim Woodward (2000). To say that a generalization is change-relating is

to say that it describes a relationship between two or more variables in which an intervention that changes one variable will bring about a change in another variable. For instance, the Boyle-Charles law can be regarded as a change-relating generalization, since intervention on one variable, say the temperature of an enclosed volume of gas, can bring about a change in another variable, the pressure of the gas.<sup>1</sup> The stipulation that these generalizations be *direct* is meant to eliminate cases where a change in one part brings about a change in another part via the action of intervening parts.

Perhaps the most notable difference between the complex-systems and Salmon/Railton approach is that Salmon/Railton mechanisms are *sequences of interconnected events* while complex-systems mechanisms are *things* (or objects). When Salmon or Railton talk about mechanisms, they are generally talking about a chain or web of events leading up to a particular event. For instance, they might speak about the chain of events that lead to the breaking of a certain window: a boy hit a baseball; the baseball ricocheted off the tree and crashed into the window. While the sequence of events leading to the breaking of the window certainly involves some entities that are stable enough to be called objects—the boy, the baseball, the bat, the tree, and the window—the complex of these objects do not form a stable enough configuration to be called an object.

Compare these examples to examples of mechanisms on the complex-systems conception: for instance, watches, cells, organisms, and social groups. These mechanisms are systems consisting of stable arrangements of parts. In virtue of these arrangements, the systems as a whole have stable dispositions—the behaviors of these mechanisms. These dispositions can manifest themselves at more than one time and place. In this sense, the behavior of a complex-systems mechanism is general. The watch, for instance, continues its periodic motion more-or-less indefinitely. Complex-systems mechanisms are general in a second sense. Although any particular mechanism will occupy a particular region of space-time, it is an important feature of our world that it often contains many tokens of a single type of mechanism. For instance, the human central nervous system contains around a trillion neurons. There are lots of human beings, as well as lots of other organisms, that have neurons whose structure is similar to that of human neurons. Consequently, one can develop a general model of neurons that subsumes countless neural events.

1. In earlier papers (Glennan 1992, 1996) I have called generalizations describing interactions between parts “laws.” I did so with the caveat that these laws must be understood in a more homely way than philosophers typically understand them. Woodward and others have convinced me that, given that many philosophers think laws must be exceptionless, my use of the term “law” was liable to lead to misunderstanding. For more on the relationship between laws and change-relating generalizations, see the online version of this paper.

A second difference between the Salmon/Railton and complex-systems conceptions of mechanism has to do with the treatment of interactions. At a general level, there is a similarity between the accounts. The parts referred to in (M) are objects, and objects can be understood as causal processes in Salmon's sense. Interactions between these parts will then be intersections in causal processes that introduce changes to the persistent structure of those processes, i.e., changes to the properties of the parts. However, there is a decided difference between the explication of interactions. While I have characterized interactions using direct invariant change-relating generalizations, Salmon has recently (1994) characterized them in terms of exchange of conserved quantities. For instance, two colliding particles interact when they exchange a conserved quantity like momentum. Salmon moved to the conserved quantity approach to avoid certain objections to the counterfactual approach (Kitcher 1989), and while the move has the advantage of characterizing interactions in terms of physical theory rather than the semantics of counterfactuals, it has the disadvantage that it obscures similarities between kinds of interactions among higher-level entities. Consider, for instance, a social mechanism whereby information is disseminated through a phone-calling chain. The people in the chain are the mechanism's parts and each phone call represents an interaction between parts of the system. Now each such interaction must, in virtue of the supervenience of everything upon physics, involve a string of physical events in which conserved quantities are exchanged. However, to say that my phone call involves the exchange of a conserved quantity is utterly uninformative about the nature of particular interactions over the phone. To talk in terms of exchange of conserved quantities requires us to treat mechanisms at a level at which such talk is intelligible, namely at the level of physics. But different tokens of a single higher level event type (e.g., a phone call event type) may have nothing in common in terms of their micro-physical descriptions. Thus, if interactions are characterized in terms of exchange of conserved quantities, tokens of higher level interactions cannot be recognized as forming higher-level kinds.

**4. Mechanistic Explanation on the Complex-Systems Approach.** Let us consider how to develop an account of mechanistic explanation consistent with the complex-systems approach. Philosophers have generally agreed that scientific explanations come in (at least) two varieties—explanation of singular events and explanation of general regularities. Let us begin with the mechanistic explanation of regularities. Roughly, the idea is this: To mechanistically explain a regularity, one describes a mechanism whose behavior is characterized by that regularity. For instance, a mechanistic explanation of Mendel's law of segregation will describe the meiotic mech-

anism that produce gametes and show how this mechanism creates (assuming certain background conditions) an equal number of gametes containing each allele of a given locus.

To spell this account out in more detail, we need to introduce the notion of a mechanical model:

- (MM) A mechanical model is a description of a mechanism, including (i) a description of the mechanism's behavior; and (ii) a description of the mechanism which accounts for that behavior. (Glennan under review)

The two-part characterization of a mechanism, in terms of a behavior and the mechanism that produces it, leads naturally to a two-part characterization of a mechanical model. The behavioral description is a description of the external behavior of a mechanism. The mechanical description is a description of the internal structure—the guts of the mechanism. The distinction between behavioral and mechanical descriptions is roughly the distinction between what a system is doing and how it is doing it.

For the purpose of illustrating how mechanisms can explain regularities, we will consider the behavior of what I call an input-output mechanism. An input-output mechanism is a complex system that is situated in its environment in such a way that there are characteristic environmental events (inputs) that trigger a sequence—perhaps multi-stranded—of interactions between parts of the mechanism. This sequence concludes with some terminating event, the output. A familiar example of an input-output mechanism is a computer running a word processing program: key presses are inputs while the characters that appear on the screen are outputs.

The behavior of input-output mechanisms can be characterized by the same type of invariant change-relating generalizations that characterize interactions between parts. Take such a generalization as an explanandum. Evidently, to the extent that a description of the internal workings of the mechanism elucidates the process by which the parts of the mechanism, triggered by the input event, interact in order to produce the output event, that description can be construed as the explanans. To put it more succinctly, the mechanical description associated with the mechanical model illustrates why the behavioral description is true.

The process of mechanistic explanation requires one to formulate, perhaps in very sketchy terms, a mechanical model. The two parts of this model, the behavioral and the mechanical descriptions, are respectively the explanandum and explanans in the mechanistic explanation. Because these descriptions are generally statements, the mechanical model can be construed as an argument in which the statements in the mechanical description are premises and the conjunction of statements in the behavioral

description is the conclusion. So construed, explanations apparently are arguments, just as Hempel and Kitcher suggest. I would like to resist this conclusion. While it is sometimes the case that a description of the parts of the mechanism will entail a description of the mechanism's outward behavior, the explanation lies not in the logical relationship between these descriptions but in the causal relationships between the parts of the mechanism that produce the behavior described.

The mechanisms so far discussed are systems consisting of relatively stable configurations of parts that give rise to robust behaviors which can be expressed by invariant generalizations. I have argued elsewhere (Glennan 1996, 1997) that most of the generalizations that scientists call laws are in fact descriptions of the behavior of mechanisms. Mendel's laws describe aspects of the behavior of mechanisms that transmit genetic material; Kepler's laws describe aspects of the behavior of gravitational mechanisms, and so on. Laws of this kind I call mechanically explicable laws. While most laws are mechanically explicable, inevitably there must be some laws that are not. For instance, given our present understanding of electricity and magnetism, it appears that Maxwell's equations are not mechanically explicable. There is not, for instance, a mechanical ether consisting of particles whose interactions could explain the propagation of electromagnetic waves. Laws such as these, which I call fundamental laws, represent brute nomological facts of our universe.

Mechanisms consisting of complex systems with stable relationships between parts can explain many of the regularities that occur in the world. But what of the explanation of particular events? There are two different sorts of events that we might seek to explain. On the one hand, there are events that are the product of the operation of reliable mechanisms. On the other, there are events that are not the product of such mechanisms. These latter events I call "genuinely singular events."

Explanation of events that are the product of reliable mechanisms follow a pattern familiar from standard accounts of explanation by nomic subsumption. Consider, for instance, the explanation of the fact that my son John has blue eyes. Blue eyes involve expression of a recessive gene. Given that my wife has blue eyes and I do not, she must carry two copies of the recessive allele while I carry one. Consequently, given the mechanisms of gamete formation, reproduction, and expression of genes determining eye color, there is a probability of 50% (relative to a generally satisfied set of background conditions) that John would have blue eyes, as in fact he does. Our explanation looks like this:

My wife is homozygous for the recessive allele for blue eyes.  
I am heterozygous for the recessive allele for blue eyes.  
John is our biological child.

If a child has one parent who is homozygous for a recessive allele and one who is heterozygous, the probability of the child being homozygous is 50%

---

John has blue eyes. [p = .5]

This is a typical example of explanation by nomic subsumption (except that it would not meet Hempel's high probability requirement). Its only connection with mechanistic explanation is that the statistical "law" which serves as its fourth premise is mechanically explicable.

While arguments of this kind are explanatory, it is also possible to ask for a different kind of explanation of the event in question. Rather than show how a certain mechanically explicable law gives a certain event some degree of probability, it is possible to explain this event by going through a play-by-play description of the operation of the relevant mechanism on that occasion. In this case, one could tell the story of the formation of the egg and sperm, the events leading to the particular sperm fertilizing the particular egg, and the subsequent developmental sequence leading to the expression of the particular gene combination in John's blue eyes. While explanations of this kind are possible, it is generally preferable with events of this sort to pursue a two-part explanatory strategy. First, provide a nomic subsumption explanation of the event, and, second, provide a mechanistic explanation of the law or invariant regularity under which the event is subsumed. This strategy explains both what caused the particular event in question and how that cause fits into a common pattern.

Kim Sterelny (1996) has offered some useful terms to characterize the difference between the two types of explanations of a particular event like my son's developing blue eyes. Sterelny suggests that, depending on our explanatory interests, we can explain a particular event either as the result of a robust process or as the result of an actual sequence. To illustrate this claim, he considers two different explanations of the outbreak of World War I (1996, 195). The actual-sequence explanation describes in detail the events leading to declarations of war in August 1914—the visit of the Archduke to Sarajevo, his assassination, the Austrian response, and so on. The robust-process explanation offers a more general description of the system of alliances, mobilization plans, nationalism, militarism and other features of Europe in 1914 that in retrospect suggest that World War I was more or less bound to happen. Sterelny argues that these explanations are both correct, but are appropriate answers to different questions.

Sterelny appears to believe that one can explain any event either from the robust process or actual sequence perspective. However, if, as I have suggested, there is a class of genuinely singular events that are not products of the operation of reliable mechanisms, such events could not be given



robust process explanations. Consider an explanation of how I met my wife. That event, like many events in which people meet, involved a confluence of events that were not to be expected and will not be repeated. In our case, for instance, an explanation of the event would, at the very least, involve an explanation of why each of us chose to attend graduate school at the University of Chicago, an explanation of the factors that led us to apply for student housing, and an explanation of the process that led the housing office to place each of us in the same building. This is an actual sequence explanation *par excellence*, but there does not appear to be a corresponding robust process.

Given that robust processes arise from the operation of reliable mechanisms, it might seem that events that are not the product of robust processes cannot be given a mechanistic explanation. Here, however, it is useful to recall the distinction made above between mechanisms as things and mechanisms as processes. The actual sequence of events is a process; it is simply not a *robust* process. Mechanisms in the sense of Salmon and Railton can characterize the actual sequence of events and connecting processes that bring about a particular explanandum event, even if that sequence does not represent the operation of a robust process. Such an explanation can be called a fragile actual-sequence explanation.

Even though mechanisms as complex systems cannot explain fragile actual sequences, there is still one point where the Salmon/Railton account can be improved upon by reference to the complex-systems conception of mechanism. This point concerns how an “actual sequence” of events is specified. Because events in the universe are capable of all sorts of descriptions at various levels of spatial and temporal grain, it is not exactly clear what could be meant by the “actual sequence of events.” In the case of World War I, what events are in this actual sequence? Sterelny suggests that this sequence includes such things as the firing of a pistol at the Archduke, but why do we describe events at this level? Why not, for instance, describe the trajectories of molecules in the Archduke’s body?

The reason for this choice is that, even in fragile actual-sequence explanations, the description of a process is guided by considerations of robustness and reliability. The “actual sequence” of events leading up to the beginning of the First World War could aptly be described as a *fragile process with robust parts*. This description is apt because, even in actual-sequence explanations, we characterize processes in terms of interactions of robust entities. Gavril Princip, the Archduke, his driver, etc. are all robust entities whose interactions comprise the sequence of events leading to the First World War. The particular arrangement of the parts (e.g., the location of the Archduke and the Assassin on the fateful day) may be ephemeral, the parts themselves (the Archduke and the Assassin) are not. Consequently, it was predictable what was going to happen when one part,

the Assassin, used another part, the Gun, at close range to shoot at a third part, the Archduke. Much the same thing can be said about my earlier example of the ball breaking the window. While the particular configuration of ball, bat, tree and window is not robust, those things themselves are robust. In both cases, it is the robustness of these parts that accounts for our decision to describe actual sequences in terms of interactions between them.

**5. Unification by Common Mechanisms.** Whatever the merits of their particular proposals, advocates of unification models are right to insist that scientific explanations generally unify our understanding of the world. How can the proponent of causal-mechanical explanation account for this property of explanation? Salmon makes the following proposal:

We explain events by showing how they fit into the causal nexus. Since there seem to be a small number of fundamental causal mechanisms, and some extremely comprehensive laws that govern them, the ontic conception has as much right as the epistemic conception to take the unification of natural phenomena as a basic aspect of our comprehension of the world. The unity lies in the pervasiveness of the underlying mechanisms upon which we depend for explanation. (Salmon 1984, 276)

In concluding this paper I will explore whether Salmon's proposal does justice to the idea of explanation as unification.

When Salmon speaks of "fundamental causal mechanisms" governed by "extremely comprehensive laws," he seems to have in mind some set of basic types of microphysical interactions, governed by what I have termed fundamental physical laws. Salmon's view is that mechanistic explanations provide unifying understanding of the world by appealing to these laws. This view is consistent with some of the best examples of explanatory unification. Consider, for example, Newton's explanation of the motion of the planets, the behavior of projectiles, and the changing tides. In showing that all of these phenomena can be explained in terms of the operation of a single fundamental force, Newton engaged in an impressive act of explanatory unification.

While there are other examples in which we appeal to fundamental laws to explain a variety of apparently diverse phenomena, there are many kinds of unification that do not involve appeal to *fundamental* mechanisms or laws. Our discussion in the preceding section showed how characterization of the mechanisms that give rise to robust processes often takes one away from the microphysical detail of actual sequences. Robust higher-level mechanisms are, to use Wimsatt's (1994) phrase, dynamically

autonomous—that is, the higher level processes are relatively insensitive to changes in the microphysical processes on which they supervene.

The complex-systems approach to mechanisms does not suppose that unification derives from unity of fundamental mechanisms. According to the complex-systems approach, mechanisms are collections of parts and parts are objects, but the objects that are parts of mechanisms may themselves be complex structures. An information processing mechanism used in human cognition may have interacting parts in the sense required by the complex-systems conception, even if these parts are spatially distributed and realized by different neurological components in different instances of the mechanism. This and other examples show that the explanatory unification afforded by the mechanistic approach derives not only from the commonality of fundamental laws but from the existence of mechanisms that have a common higher-level structure even if they differ in microstructure.

Even if the complex-systems approach can provide a reasonable account of how mechanistic explanations unify, there remains an important metaphysical controversy over the priority of the concepts of explanation and causation. The unificationist argues that explanations precede causes while the mechanist argues that mechanisms precede causes. Kitcher represents the unificationist position as follows:

The heart of the unification approach is that we cannot make sense of the notion of a basic mechanism apart from a systemization of the world in which as many consequences as possible are traced to the action of as small a number of basic mechanisms as possible. (Kitcher 1989, 497)

The mechanist has a simple answer to this: Kitcher has confused epistemological and metaphysical questions. The claim that we identify basic mechanisms by choosing those mechanisms that maximally unify our understanding of diverse phenomena is really a claim about how we discover and test theories about mechanisms. It is related to the claim that we choose scientific theories on the basis of their simplicity and scope. While there is undoubtedly truth in Kitcher's claim about how we discover and test theories about fundamental mechanisms, we can still "make sense of the notion of a basic mechanism" apart from this global perspective. In fact, our concept of causation suggests that events occurring at some point in space and time are explained as the consequence of the operation of causal mechanisms operating in that region of space and time. Our global evidence suggests that—quantum mechanics aside—causality is everywhere local.

This brief argument will probably not convince instrumentalists of the reality of causes, but I hope it is enough to show that Kitcher's position

is not the only possible interpretation of the global character of theory choice. And if it is possible to treat causes as prior to explanations, it is also possible to see how causal explanations can provide explanatory unification.

## REFERENCES

- Bechtel, William and Robert Richardson (1993), *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton: Princeton University Press.
- Glennan, Stuart (1996), "Mechanisms and the Nature of Causation", *Erkenntnis* 44: 49–71.
- (1997), "Capacities, Universality and Singularity", *Philosophy of Science* 64: 605–626.
- (under review), "A Model of Models".
- Kitcher, Phillip (1981), "Explanatory Unification", *Philosophy of Science* 48: 507–531.
- (1989), "Explanatory Unification and the Causal Structure of the World", in Kitcher and Salmon (eds.), *Scientific Explanation: Minnesota Studies in the Philosophy of Science*, Vol. 13. Minneapolis: University of Minnesota Press, 410–505.
- Kitcher, Phillip and Wesley Salmon (eds.) (1989), *Scientific Explanation: Minnesota Studies in the Philosophy of Science*, Vol. 13. Minneapolis: University of Minnesota Press.
- Machamer, Peter, Lindley Darden and Carl Craver (2000), "Thinking about Mechanisms", *Philosophy of Science* 67: 1–25.
- Railton, Peter (1978), "A Deductive-Nomological Model of Probabilistic Explanation", *Philosophy of Science* 45: 206–226.
- (1981), "Probability, Explanation, and Information", *Synthese* 48: 233–256.
- Salmon, Wesley (1984), *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
- (1989), "Four Decades of Scientific Explanation", in Kitcher and Salmon (eds.), *Scientific Explanation: Minnesota Studies in the Philosophy of Science*, Vol. 13. Minneapolis: University of Minnesota Press, 3–219.
- (1994), "Causality without Counterfactuals", *Philosophy of Science* 61: 297–312.
- Sterelny, Kim (1996), "Explanatory Pluralism in Evolutionary Biology", *Biology and Philosophy* 11: 193–214.
- Wimsatt, William (1994), "The Ontology of Complex Systems: Levels, Perspectives and Causal Thickets", *Canadian Journal of Philosophy* 20 (supplement): 207–274.
- Woodward, James (2000), "Explanation and Invariance in the Special Sciences", *British Journal for the Philosophy of Science* 51, 197–254.