

Mechanisms and Laws: Clarifying the Debate

Carl F. Craver and Marie I. Kaiser

To appear in: Chao, H.-K./ Chen, S.-T./ Millstein, R. L. (2013) (eds.): *Mechanism and Causality in Biology and Economics*. Dordrecht: Springer.

Abstract

Leuridan (2011) questions whether mechanisms can really replace laws at the heart of our thinking about science. In doing so, he enters a long-standing discussion about the relationship between the mechanistic structures evident in the theories of contemporary biology and the laws of nature privileged especially in traditional empiricist traditions of the philosophy of science (see e.g. Wimsatt 1974; Bechtel and Abrahamsen 2005; Bogen 2005; Darden 2006; Glennan 1996; MDC 2000; Schaffner 1993; Tabery 2003; Weber 2005). In our view, Leuridan misconstrues this discussion. His weak positive claim that mechanistic sciences appeal to generalizations is true but uninteresting. His stronger claim, that all causal claims require laws,

is unsupported by his arguments. Though we proceed by criticizing Leuridan's arguments, our greater purpose is to embellish his arguments in order to show how thinking about mechanisms enriches and transforms old philosophical debates about laws in biology and provides new insights into how generalizations afford prediction, explanation and control.

1. Introduction

Over a decade ago Machamer, Darden and Craver (2000) suggested that the philosophy of science, especially the biological sciences, could usefully be reconfigured by thinking about how scientists construct, evaluate, and revise their understanding of mechanisms. They boldly asserted that traditional philosophical topics such as causation, discovery, explanation, functions, laws, levels, models, and reduction would be fundamentally transformed by recognizing the centrality to many areas of science of the search for mechanisms. The revolution they envisioned replaced the last vestiges of the once-received positivist gestalt with a new mechanistic vision, expressed in the very language in which scientists talk about their work and sensitive to problems faced within mechanistic research programs in areas as diverse as biology, cognitive science, ecology, and neuroscience.

Though this way of thinking about the philosophy of science has gained rapid and widespread acceptance, it has unsurprisingly attracted a good deal of criticism from those who wonder whether the mechanical philosophy is really as revolutionary as its propo-

nents suggest and from those who think that traditional ways of thinking about the philosophy of science address problems that the mechanical philosophy is ill-equipped to handle. And one might be forgiven for thinking that there is no more central battle ground in that debate than the perennial issue of laws of nature. Positivist philosophy of science and its descendents place the concept of a law of nature at the very heart of their thinking about causation, explanation, prediction, and reduction in particular. From that traditional vantage point, it is reasonable to ask precisely how the concept of mechanism, which plays many of the same roles in the new paradigm, is related to the concept of a law of nature.

So conceived, one naturally sees the concept of mechanism as replacement for the concept of laws. And indeed, a casual reading of the mechanistic literature would give the impression that this is precisely what the mechanists intended to do. Mechanists regularly note that the term 'law' is descriptively out of place in the biological sciences. Biologists and other scientists of the middle range (neuroscientists, physiologists, psychologists, etc.) seem to avoid the term 'law,' and conceive of their work instead in terms of the discovery

of mechanisms. Furthermore, the mechanist's rejection of a law-centered picture of science is a part of their general rejection of the "Euclidean ideal" (Schaffner 2008) of science, according to which knowledge is arranged in closed deductive axiomatic systems with strict law statements as the axioms. How, they ask, would the philosophy of science look if this formal gestalt, which had already worn quite thin in places, were replaced by a more material, mechanistic, gestalt: one emphasizing the causal structures that scientists much more frequently discuss (see Craver 2002)? Seeing the mechanistic project in this light leads one to ask, as Bert Leuridan (2010) does in a recent paper in *Philosophy of Science*, whether mechanisms can really replace laws at the heart of our thinking about science.

Leuridan believes they cannot. In this paper, we assess his arguments. Though his arguments, as we show below, leave one with no compelling reason to maintain the traditional view, his discussion demonstrates the need for greater clarity about the place of laws in mechanistic sciences.

But first the ground rule: All parties to this discussion, as Leuridan points out, agree that the traditional notion of a “strict law,” the universally quantified material conditional with unrestricted scope and a good deal besides, has little application in biology and other special sciences. Mechanists have openly embraced a number of arguments for this conclusion, most notably John Beatty’s (1995) suggestion that the laws of biology are evolutionarily contingent and Stuart Glennan’s (1996, 2002) idea that the generalizations of biology are mechanistically fragile, and so probabilistic and prone to breakdown. Other mechanists emphasize that theoretical claims in biology are typically limited in scope, applying only to some species and strains (cf. Hull 1978), and that the scope of such generalizations is restricted to life on earth in a particular epoch (cf. Smart 1963). Whatever the reason, mechanists have been happy to echo these criticisms as evidence of the limited applicability of the traditional law-based view to the philosophy of biology. But contra Leuridan’s suggestion, it should also be noted that the mechanist’s general opposition to strict laws does *not* entail opposition to the idea that biologists and other scientists of the middle range seek to

learn about and describe general facts. None of these arguments showing that the idea of a strict law distorts crucial features of biology shows that there are no general facts about biology or that generalizations play no important role in biological research practice. And no mechanist has ever made such claims.

Though we proceed by criticizing Leuridan's arguments, we have a larger purpose, namely, to illustrate how thinking about mechanisms enriches and transforms the philosophical debate about the role of laws in biology. In our view, the debate over whether or not there are laws in biology has outlived its usefulness. Nobody anymore denies that there are stable regularities that afford prediction, explanation, and control of biological phenomena. Whether such stable regularities count as laws depends on what one requires of laws, but it is undeniable that generalizations of this sort do many kinds of work in biology. What remains is the admittedly difficult work of showing how this is possible. If one takes the biological sciences to be largely dedicated to the search for mechanisms, in contrast, one can begin to ask in relatively precise ways how generalization contributes to the search for mechanisms and, conversely, what

the idea of mechanism brings to long-standing questions about how generalizations afford prediction, explanation and control.

We begin by clarifying Leuridan's thesis and his central ontological and epistemological arguments (Section 2). In Section 3 we consider Leuridan's ontological claims and argue that Leuridan fails to show that mechanisms must involve regularities (Section 3.1 and 3.2) or that there must be fundamental laws without underlying mechanisms (Section 3.3). Despite the emphasis Leuridan places on the notion of projection (i.e., extrapolation), he fails to explain why the generalizations of biology are stable and why certain facts can be extrapolated while others cannot (Section 3.1). We show further how the mechanistic perspective provides new resources to ameliorate these extrapolation problems. In Section 4 we turn to the epistemological issues. We reject Leuridan's claim that mechanistic models must contain law statements, and we show how mechanistic knowledge contributes to the search for stable generalizations. We conclude that continued debates over whether mechanisms can replace generalizations are likely to be unproductive. We conclude, second, that by taking a mechanistic stance, one gains a new vantage

point on old problems about laws and a view to new problems about the construction, evaluation, and revision of models of biological mechanisms.

2. Leuridan's Thesis

In his title, Leuridan asks, "Can mechanisms really replace laws of nature?" He answers, "No". In fact, Leuridan's positive thesis is much weaker than this title suggests.

Before formulating this weaker claim, it is necessary first to clear up some terminology. Leuridan defines laws as "generalization[s] describing a regularity, not some metaphysical entity that produces or is responsible for that regularity" (2010, fn 1). This definition ignores three traditional distinctions that have brought much-needed clarity to the discussions of laws in the philosophy science. First, we distinguish laws (metaphysical entities that produce or are responsible for regularities) and law-statements (descriptions of laws). If one does not respect this distinction, one runs the risk (as Leuridan does) of unintentionally suggesting that sentences, equations, or models are responsible for the fact that certain stable regu-

larities hold. In like fashion, we distinguish regularities, which are statistical patterns of dependence and independence among magnitudes, from generalizations, which describe regularities. Finally, we distinguish regularities from laws, which produce or are otherwise explain the patterns of dependence and independence among magnitudes (or so one might hold).

Let us now reconstruct Leuridan's real thesis. First, Leuridan endorses the ground-rule of our discussion. Strict law statements, as Leuridan understands them, are nonvacuous, universally quantified, and exceptionless statements that are unlimited in scope, apply in all times and places, and contain only purely qualitative predicates (2010, 318). Noting that few law-statements in any science live up to these standards, Leuridan argues that the focus on strict law statements (and presumably also on strict laws) is unhelpful for understanding science. Instead, he focuses on the concept of a pragmatic law (or p-law). Following Sandra Mitchell (1997, 2000, 2003, 2009), Leuridan understands p-law statements as descriptions of stable and strong regularities that that can be used to predict, explain, and manipulate phenomena. A regularity is stable in proportion to

the range of conditions under which it continues to hold and to the size of the space-time region in which it holds (2010, 325). A regularity is strong if it is deterministic or frequent. P-law statements need not satisfy the criteria for strict law statements. Thus, Leuridan's question is not whether mechanisms can replace laws, simpliciter. Rather it is whether mechanisms can replace p-laws and, correlatively, whether descriptions of mechanisms can replace p-law statements in our thinking about science.

Yet Leuridan's thesis is narrower still. He distinguishes two "kinds" of mechanism: complex system mechanisms (cs-mechanisms) and Salmon/Railton mechanisms. Leuridan characterizes cs-mechanisms as stable configurations of robust objects that produce stable behaviors (2010, 319; see also Glennan 2002, 344-46).¹ Leuridan does not define Salmon/Railton mechanisms, except

¹ The full passage is: "Contrary to Salmon/Dowe mechanisms, complex systems mechanisms (cs-mechanisms) are robust and stable. They form stable configurations of robust objects, and as a

to say that they involve causal processes and causal interactions (2010, 319). However, if we follow Glennan, they might be understood as “sequences of interconnected events” or “a chain or web of events leading to a particular event” such as: “a boy hit a baseball; the baseball ricocheted off the tree and crashed into the window” (Glennan 2002, 345). Salmon/Railton mechanisms are singular causal chains. Substituting into Leuridan’s title question yields something closer to the question he in fact addresses: “Can cs-mechanisms really replace p-laws?” Leuridan is not always clear to distinguish this ontological question from its epistemological twin: “Can models of cs-mechanisms replace p-law statements?”² But in either case, he concludes they cannot.

whole they have stable dispositions: the overall behaviors of these mechanisms.” (Leuridan 2010, 319)

² On closer inspection, Leuridan’s question is still too imprecise since it does not specify the purpose for which cs-mechanisms are intended to replace p-laws (or vice versa). Mechanisms and stable generalizations serve many functions in our thinking about science.

More precisely, Leuridan presents four theses, two of which he describes as ontological, and two of which he describes as epistemological:

First, [cs-] mechanisms are ontologically dependent on stable regularities. There are no [cs-]mechanisms without both macrolevel and microlevel stable regularities. [L1]

Second, there may be stable regularities without any underlying [cs-]mechanism. [L2]

Third, models of [cs-]mechanisms are epistemologically dependent on pragmatic laws. To adequately model a [cs-]mechanism, one has to incorporate pragmatic laws. [L3].

Finally, pragmatic laws are themselves not epistemologically dependent on mechanistic models. They need not always refer to a mechanism underlying the regularity at hand. [L4]

(2010, 318-19)

Perhaps mechanisms are useful for some philosophical purposes and laws are useful for others.

We have inserted the qualification to cs-mechanisms specifically, given that Leuridan offers no argument even purporting to show that non-cs-mechanisms are dependent upon p-laws (see Bogen 2005, 2008), a gap to which we return below. Leuridan argues that cs-mechanisms cannot replace p-laws in our thinking about the ontology of science because cs-mechanisms are ontologically dependent on p-laws (L1) but the opposite is not the case (L2). Furthermore, he claims that cs-mechanisms cannot replace p-laws in our thinking about epistemology because models of cs-mechanisms are epistemically dependent on p-laws (L3) and not *vice versa* (L4).

Before we address Leuridan's arguments, it is necessary first to set the record straight. When Leuridan asks "Can cs-mechanisms really replace p-laws?" the word "really" suggests that somebody has claimed that they can. Is this true? Do mechanists really insist that scientists can discover, explain, predict, and control the action potential, heredity, long-term potentiation, natural selection, and neurotransmitter release (to name a just few of the lengthy examples that mechanists have discussed) without forming generalizations about them? Do mechanists think that the Hodgkin-Huxley model of

the action potential, the theory of evolution by natural selection, and the current models of long-term potentiation and neurotransmitter release make no appeal to regular occurrences? In a word, no. They are quite explicit on this matter (see for instance Bechtel and Richardson 2010, 232; Glennan 1996, 52; 2002, 345; Machamer, Darden and Craver, 2000, 3, 7; Bechtel and Abrahamsen 2005, fn. 1, 437; Craver 2007, Ch. 3, 233-34). James Bogen (2005), the mechanist most critical the role of generalizations and regularities in our thinking about causation, stresses at great length the importance of Mitchell's treatment of p-law statements (and the regularities they describe). He also emphasizes the many epistemic roles that generalizations play in the search for mechanisms (cf. Bogen 2005, 401):

- (a) to describe the phenomenon to be explained;
- (b) to suggest and sharpen questions about causal mechanisms;
- (c) to describe constraints on acceptable mechanistic models;
- (d) to measure or calculate quantities relevant to the mechanism;
- (e) to support inductive inferences without which mechanisms could not successfully be studied;

- (f) to support extrapolation of mechanistic knowledge to new cases;
- (g) to design effective experiments to test mechanisms;
- (h) to simulate the behavior of mechanisms.

The list could no doubt go on. In short, no mechanist denies that biologists search for regularities and routinely formulate generalizations (p-law statements) that can be used for prediction, explanation, and control of phenomena. Indeed, it is hard to see how any significant human activity could be pursued without discovering and representing (in some sense) such regularities. The mechanist claims simply that it is useful to ask further about the material structures that generalizations describe, and about how this affects the various tasks scientists perform. In many areas of science, scientists seek to describe mechanisms in order to explain, predict, and control phenomena. If one places the idea of mechanism at the center of one's thinking about those sciences, one suddenly sees p-laws in a new light, with new roles to play (compare Bogen's list to Leuridan's emphasis on prediction, explanation, and control). The question is not whether biological phenomena operate in accordance with p-

laws or exhibit p-regularities, but rather how the search for those regularities fits into the central aim of describing mechanisms.

3. The Ontology of Mechanisms

Let us now consider Leuridan's argument that p-laws are ontologically fundamental to cs-mechanisms. He argues for two component theses (2010, 329):

- (a) There can be no cs-mechanism without some stable behavior produced by that mechanism (Leuridan calls this the "macrolevel regularity"); and
- (b) There can be no cs-mechanism without some regular behaviors, operations, or activities displayed by or engaged in by the mechanism's component parts (Leuridan calls these "microlevel regularities").

Given Leuridan's definition of a cs-mechanism (cf. 2010, 319; see also Section 2), (a) is a tautology: "There can be no stable configurations of robust objects that produce stable behaviors without some stable behavior produced by that mechanism." We agree.

Surprisingly, Leuridan offers historical evidence to shore up his case. “In the life sciences,” he writes, “reference to mechanisms cannot be detached from matters of projectability.” (2010, 329) For example, he notes, Thomas H. Morgan intended his work on the mechanisms of heredity in *Drosophila* (see Morgan et al. 1915), work summarized in a book aptly titled *The Mechanisms of Mendelian Heredity*, to apply outside of the laboratory and to other organisms as well.

Mechanists deny neither that Morgan sought projectable generalizations nor that he succeeded in finding them. However, a mechanist might well insist that scientists sometimes seek details about a particular causal mechanism without any interest in generalizing to other cases. Evolutionary biologists might describe the mechanisms that increased the prevalence of a single adaptive trait in a population or that produced a single speciation event. Epidemiologists might be interested in how AIDS first came to South Korea. Archeologists might be interested in the origins of maize cultivation in North America. Ecologists might be interested in the mechanisms causing fish populations to dwindle in the Chesapeake Bay. Morgan

was interested in generality, we grant, but sometimes scientists just want to know how some particular event came to pass, and so they describe the particular mechanism that is responsible.

Bogen (2005, 2008) argues persuasively that the concept of causation (understood in terms of “causally productive activities” (2008, 112)) and the concept of regularity can (and do) come apart from one another.³ We have no difficulty imagining quite irregular mechanisms, such as the mechanisms of neurotransmitter release, that work roughly ten percent of the time, or a rusty chainsaw that starts arbitrarily infrequently. What matters for the existence of a mechanism is not how frequently it runs in the limit, but how it works when it works, however infrequently. Viewed from this perspective, singular, unrepeatable causal chains (so-called one-off

³ Leuridan mentions Bogen’s work, but rejects it summarily on the ground that Bogen’s criticism of regularism relies on a strict (i.e., universal) notion of ‘regularity’ (see 2010, 330). But this is false. Bogen frames his discussion in terms of Mitchell’s view of generalizations.

mechanisms or Salmon/Railton mechanisms) are a special, limiting case of *cs*-mechanisms, not something altogether different. While Leuridan's thesis that there can be no *cs-mechanisms* without some stable behavior produced by that mechanism (cf. 2010, 330) is tautologically true, Leuridan's unqualified thesis that "there are no *mechanisms* without... macrolevel... stable regularities" (2010, 318; our emphasis) is clearly false. One-off mechanisms are mechanisms *without* a macrolevel regularity. So much for the ontological claim.

3.1 *Extrapolation of Generalizations*

Things look a bit more promising if we reconstruct Leuridan's projectability thesis as a purely *epistemic* thesis. Morgan wanted to apply what he learned about the mechanisms of heredity by studying *Drosophila* in the lab both to flies outside the lab and to other species. Surely the mechanist owes some kind of story about how this is possible. The clear solution, one might think, is to recognize that there are laws – however exception-ridden, probabilistic, and mechanistically fragile – that license this application. And one might insist that Morgan referred to, and indeed formulated, Men-

del's second law while making a career of discovering exceptions to independent assortment (see Allen 1978; Darden 1991). Scientists form generalizations, and then they use those generalizations to say what will happen in new cases. Of course, no mechanist denies that induction and extrapolation (or projection) are important to science. But how are p-laws supposed to help with this task? If p-laws are merely law *statements*, as Leuridan defines them, then they are clearly not the kind of thing that can explain why a given regularity is stable and strong. Law statements *express that*, but do not *explain why*, certain regularities are stable and strong. It seems we must understand Leuridan to mean that stable p-regularities themselves (rather than descriptions of p-regularities) are necessary for one to extrapolate mechanistic knowledge. Here, in full, is Leuridan's discussion of stability: "What are the conditions on which the regularity under study is contingent? How spatiotemporally stable are these conditions? And what is the relationship between the regularity and its conditions (is it deterministic, probabilistic, etc.?)" (2010, 325). Given that stability is defined in terms of the range of circumstances in which a generalization holds, the epistemological thesis

that extrapolation to conditions outside of the laboratory and to conditions in other organisms requires p-regularities, again, amounts to a tautology: if the regularities discovered about *Drosophila* in the laboratory are to hold outside of the laboratory and for other organisms, then there must be organisms outside of the laboratory for which the regularity holds. If a regularity holds only in Morgan's laboratory or only for *Drosophila*, then there is nothing outside of the laboratory or in other organisms about which to extrapolate. But this is not an explanation of why knowledge extrapolates beyond the laboratory; it is simply a claim that it does extrapolate outside the laboratory. Put this point another way: by helping himself to the idea of p-laws, which are by definition stable regularities, Leuridan does nothing to *explain why* certain facts can be extrapolated and others cannot. Nor does he tell us how to discern which features of a system can be extrapolated from those that cannot. Rather, by invoking the idea of a p-law, he merely asserts *that* there is a distinction between knowledge that can be extrapolated and knowledge that cannot be extrapolated. But the bald statement that there is a difference between the predicates that project and those that do not, conditions

that project and those that do not, and times when the consequent really ought to obtain and times when it shouldn't, is not a *victory* for p-laws, but simply an *assertion* that a central problem for any theory of p-laws has a solution.

Let's push a bit deeper. For a defender of strict laws, which by definition apply always, without exception, and without limitation of scope, it is reasonably clear how knowledge of the laws would warrant extrapolation. For a defender of a robust metaphysical notion of a law, where a law is part of the structure of the world that explains (rather than merely describes) the p-regularities we observe, then knowledge of the laws would presumably warrant extrapolation. But Leuridan weakens the notion of a law so that p-laws are mere regularities and p-law statements are descriptions of these regularities; further, such descriptions are non-universal, have exceptions, and apply only in restricted regions of space-time. In effect, he turns p-laws into imperfect regularities with no robust metaphysical backing. Whether such a weakened p-law warrants extrapolation outside of the laboratory depends upon whether one in fact finds that the regularity continues to hold outside of the labora-

tory, whether the necessary background conditions hold, whether the target instance under consideration is one of the exceptions or whether it is not. P-laws, as Leuridan understands them, might not warrant extrapolation. The laws might hold only in Morgan's laboratory, after all. At the very least, if one believes that p-laws offer a solution to the problem of extrapolation, then one owes a further story about how one knows when the conditions for extrapolating the regularity have been met. Leuridan offers no such story.

Bechtel and Abrahamsen (2005), whom Leuridan picks out for particular criticism on this matter, argue on independent grounds that it is philosophically unfruitful to think about the problem of generalization (or extrapolation, in Leuridan's vocabulary)⁴ in terms of laws. In a section of their paper called "Generalizing without Laws" they criticize law-based views of generalization and develop

⁴ One might distinguish generalization (i.e. expanding the scope of the schema within a species/class) from extrapolation (i.e. expanding the scope of the schema beyond the species/class). Given that the parties to this dispute do not draw this distinction we treat them as synonymous.

an alternative, prototype-based account. Because Leuridan does not mention these arguments, we repeat them here. They argue that if one thinks of biologists as attempting to build law statements, paradigmatically represented in terms of material conditionals, then it is difficult to understand the prototypical structure of biological theories. One is tempted to think of biologists as constructing, for example, a law-statement of heredity (such as Mendel's laws). When one encounters variation in that mechanism (as Morgan did), one is tempted to package the variation into the antecedent of the conditional. In fact, however, one finds that biologists typically characterize a mechanism in a particular strain of a particular species (such as wild-type *Drosophila*) and then recognize that there will be subtle variations on that mechanism in other strains, mutants, and species. They are not looking for general law statements that cover all of them, but rather for sets of prototypical models that stand in family resemblance relations to one another (cf. Schaffner 1993). To push Bechtel and Abrahamsen's point one step further: Prototype models need not be general descriptions. Bechtel and Abrahamsen also call their view an "exemplar" (2005, 438) account, noting that models of

mechanisms often describe a particular, exemplary case. (Open a biology textbook and look at some diagrams of mechanisms; more often than not, they are cartoons of a single representative mechanism.) On such an exemplar view, generalization is extrinsic to the mechanistic models (exemplars, prototypes); that is, the model need not contain general statements or general representations at all.⁵

Leuridan's insistence that the model must contain such things is simply the imposition of a philosophical prejudice onto actual scientific models that have the capacity rather to surprise us if only we open our eyes to them. The generality of such a mechanistic model is a matter of its scope of application and not something that must be represented within the model itself. If one attempts to put the generality in the model itself, to return to Bechtel and Abrahamsen's point, the model has difficulty accommodating the variability characteristic of biological mechanisms.

⁵ This view fits with the semantic view of theories that Bechtel and Richardson embrace (cf. 2010, 232).

Curiously, Leuridan fails to consider the possibility that most p-regularities are stable and strong because they are produced or maintained by mechanisms (see, for example, Bechtel 2009; Craver 2007; Darden and Craver 2002; Glennan 2010; Steel 2008; Wimsatt 1998). Why might Morgan have expected the apparent exceptions to Mendelian heredity he discovered in his lab to apply outside of the lab and in other organisms? The simple answer is this: He expected the mechanisms of heredity outside the lab and in other organisms to be more or less similar to the hereditary mechanisms at work in his *Drosophila*. The p-laws of heredity are stable and strong precisely because there is an underlying mechanism (involving, e.g., crossing over and replication of chromosomes) that explains them.

In his book *Across the Boundaries* (2008) Daniel Steel builds on early suggestions by Darden and Craver (2002) to develop an elaborate analysis of how one can extrapolate scientific knowledge based on an understanding of the relevant mechanisms. The idea behind his “comparative process tracing” (Steel 2008, 85) approach is simple and helpful: First, one uses a variety of strategies to learn about the mechanism in the model organism. Second, one compares

the mechanism in the model organism to the mechanism of the extrapolational target at certain key junctures. That is, one compares the two mechanisms at stages at which the mechanisms are most likely to differ significantly from one another. The fewer significant differences one discovers at these key points, the stronger is the basis for the extrapolation. Crucially, one need not compare all of the entities, activities, and organizational features of a mechanism to those in the target organism in order to assess the likelihood that one's extrapolation will work: one might, for example, compare downstream (rather than upstream) portions of a mechanism, given that crucial differences downstream will indicate crucial differences earlier. Conversely, similarity at a key bottle-neck point in the mechanism might allow one to neglect any differences upstream in the mechanism to focus on what comes later (see Steel 2008, 90). Furthermore, if one is interested simply in gross qualitative differences, such as whether a given drug is positively relevant for a side-effect or negatively relevant for a side-effect, certain minute and highly-specific differences in the mechanisms might be less germane than the simple matter of whether there is a positive (excitatory) or

negative (inhibitory) causal or correlational relationship in the model organism. Thinking about underlying mechanisms, in short, provides new tools for assessing when our knowledge is likely to extrapolate and when extrapolation is more precarious.

Steel's strategies rely primarily on considering the mechanisms that underlie a regularity, but one might also justify extrapolation on the basis of antecedent mechanisms, such as the mechanism of natural selection. That is, one might claim that the hereditary mechanisms in *Drosophila* can be expected to apply outside of the laboratory and in other species because hereditary mechanisms are evolutionarily ancient and therefore widely conserved across the tree of life. As Bechtel (2009) argues, this mechanistic fact about the history of life warrants tentative (heuristic) extrapolation about closely related species: they might use the same mechanism, or a mechanism composed of similar entities and activities, or mechanisms with similar organizational structures. And one might expect evolutionarily ancient mechanisms to be more widely conserved, and so more fitting for extrapolation, than are relatively recent adaptations. This kind of heuristic is especially interesting in the present context given

that, according to this heuristic, a *singular* mechanism (the one-off mechanism that produced the tree of life as we now know it) warrants extrapolation of p-laws in extant species.

While we admit that these mechanistic contributions to our understanding of extrapolation solve neither the problem of induction nor Goodman's new riddle of induction (1955), we insist they nonetheless have considerably more content than the bare assertion that p-laws warrant extrapolation because they are stable and strong. Indeed, extrapolation is at least often justified by appeal to knowledge of mechanisms. In sum, it appears that our epistemic reformulation of Leuridan's argument runs into a dilemma. Either his claim is a tautology to the effect that mechanisms must be general if one is to form true generalizations about them, or it is a substantive epistemological thesis that extrapolation is possible only if there are p-laws. If the latter, then we have shown how Leuridan begs the question by presuming, rather than showing, that p-laws solve the extrapolation problem and by asserting, rather than defending, the disputed thesis that p-laws are required for extrapolation. Most importantly, however, we have reviewed some of the progress mecha-

nists have made in thinking about the problem of extrapolation. Focusing on mechanisms provides fruitful and substantive ways of thinking about how generalizations are extrapolated in scientific practice. It is unclear why Leuridan refuses the mechanist's help in addressing the extrapolation problem.

3.2 *Do cs-mechanisms require micro-regularities?*

Let us move on, then, to the second route (b) by which Leuridan argues that cs-mechanisms are ontologically dependent on stable regularities (L1). Leuridan claims: "There can be no cs-mechanism without some lower-level (c)P-regularities (i.e., the regular behaviors, operations, or activities displayed or engaged in by the mechanism's parts)." (2010, 331) A (c)P-regularity is a causal p-law, a P-law that that is "invariant under some range of interventions" (2010, 328). Leuridan argues for this thesis using a thought experiment. If the behaviors of all of the parts of the mechanism were to behave completely randomly, by which he means that they do what they do as the result of a "completely random internal process", "this would make it very unlikely to produce a macro-p-

regularity, let alone a (c)P regularity” (2010, 331).⁶ What shall we make of this argument?

Clearly, Leuridan’s thought experiment does not support the ontological conclusion that there can be no cs-mechanisms without some p-regularities among the parts. At most, it supports a probabilistic conclusion that cs-mechanisms are unlikely without p-regularities, and such an argument cannot support the negated exist-

⁶ It should be noted that Leuridan defines “irregularity“ in such a way as to effectively exclude discussion of stochastic mechanisms, mechanisms that work only infrequently or whose frequency of operation and stability in space vary over time. A mechanism that works with probability 0.000001 will count as regular on Leuridan’s account because one can write a generalization of the form $P(X) = 0.000001$. This is unfortunate as there are a number of interesting questions that one might ask about probabilistic mechanisms and mechanisms whose probability of working varies over time (as one might expect in systems that are regulated). Thanks to Jim Bogen for calling this to our attention.

tential quantifier in Leuridan's second ontological claim (b). The thesis that x is unlikely to have property F is consistent with the claim that x is F and, for non-zero probabilities, entails that x is possibly F (directly contradicting Leuridan's stated thesis). Although randomly behaving components such as those in Leuridan's example would not form a mechanism (given that the behavior of each is causally independent of the behaviors of the others), it is still possible that together they would produce a regularity, even a (c)P-regularity, of some stability and strength. Just how improbable this would be depends upon the number of variables and the number of values they might take. In order to make experimental progress in the discovery of causes and mechanisms, we regularly presume that regularities do not arise merely from chance. However, as the statistics attached to any causal experiment acknowledge, there is always some non-zero probability that the results of the experiment did arise strictly from chance. Now if macro regularities can obtain even among causally unconnected random events (as in Leuridan's example), then why would one ever suppose that it would be impossible for them to obtain among causally connected random events (what-

ever sense we can make of that notion)?⁷ It seems there is no interpretation of the idea of irregularly behaving components that sustains even the negative existential thesis entailed by Leuridan's modal claim that no (c)P-regularity can be produced by irregularly behaving components.

Perhaps what Leuridan means to claim is that very few of the mechanisms described in biology textbooks explain higher-level (c)P-regularities without appealing to regularly behaving components. If one wants to discuss the kinds of mechanism that biologists typically study, then one must acknowledge that there are true p-generalizations about the components of mechanisms. True enough. But this claim is entirely independent of the ontological thesis that cs-mechanisms depend on lower-level regularities. And no mechanist denies that there are true p-generalizations about the components of mechanisms.

⁷ Again, note that Leuridan is operating with a most unorthodox notion of 'irregularity'.

If the argument does not work for cs-mechanisms, it certainly will not work for mechanisms in general (as his title and introduction suggest). That is, it cannot establish, as Leuridan claims, that there can be no mechanisms without microlevel stable regularities. It seems one-off mechanisms (the ‘Salmon/Railton mechanisms’ discussed above) might well work without micro-level stable regularities. Such mechanisms probably wouldn’t be so scientifically interesting, and we might never know about them, but they might well exist.

Leuridan might, at this point, have entered a long debate about the regular character of causality. Perhaps he could endorse the view that the components in a mechanism can properly be said to causally interact with one another only if there exists a p-regularity relating events of one type to events of another type. If all mechanisms have interacting parts, and if there can be no interactions among parts without p-regularities, then there can be no mechanisms without p-regularities. That’s certainly an ontological thesis, and it’s one with a grand tradition. It’s also a view that some mechanists

(such as Bogen 2005, 2008; Machamer 2004; Darden 2006) explicitly challenge.

As we mentioned above, Bogen (2005, 2008) argues that causation and regularity are conceptually distinct. One set of Bogen's arguments turns on the implicit thesis that causation is local (or, in other words, intrinsic): that whether A causes B depends on facts about A, B, and their relation to one another, and does not depend on how other A-type things and B-type things behave when they interact. What matters instead is whether A and B are connected by some determinate sort of activity. One need not buy the metaphysics of activities to appreciate the intuitive pull of locality. Imagine a world composed only of two billiard-balls traveling through space-time toward one another until one day they clack together and fly back in the directions whence they came. Whether they interacted would seem not to depend on whether any other billiard balls ever meet, or on whether the same billiard balls ever meet again; neither is true in the world we are considering. The causal interaction is a fact about them and them alone (that is, an intrinsic feature of their

interaction); nobody else matters, and so no p-regularity (or any regularity) matters, to whether they interact.

A second kind of argument for the separability of regularity and causation turns on the possibility of causal relations that have no echo in the correlational structure of the world. For example, one might have a mutation that reduces the overall chance that one will get lung cancer (that is, the mutation has *negative* statistical relevance for cancer) but that, in a few unfortunate individuals, is, in fact, the trigger for lung cancer. And one might get lung cancer in virtue of having that mutation. One might smoke three packs a day (raising the chance of getting lung cancer) and in fact get lung cancer because of the mutation. The actual causal structure in such cases would appear to run counter to the regularities. To borrow a kind of example first described by Jonathan Schaffer (2000), we might imagine two neurons, A and B, synapsing on a third neuron, C. Suppose we know from experimental investigation that the probability of C's firing given A's firing alone is 0.5, that the probability of C's firing given B's firing alone is 0.5, and that the probability of spontaneous firing in C is 0. Now suppose that A, B, and C all fire. These

facts leave the causal facts under-determined. For in this situation, it might be that A caused C to fire, that B caused C to fire, or that both A and B caused C to fire. The difference between these possibilities cannot, *ex hypothesi*, depend on the regularities involved. It would seem that there is a further fact about the actual causal structure of the situation. Regularities, it might be thought, provide evidence about the causal structure of a mechanism. But the causal structure of the mechanism is something over and above the regularities by which that structure can be detected.

We do not insist on the view that causation is intrinsic, actual, and singular. We simply note that Leuridan does not address the heart of the debate about whether regularities are more fundamental than causation and mechanisms. Some philosophers, most explicitly Glennan (2002), Woodward (2002), and Craver (2007), appear to agree (to a first approximation) with the idea that the interactions in a mechanism should be characterized in terms of invariant change relating generalizations. They stress, for example, that knowledge of causes is practically valuable precisely because it is general. And they emphasize the close connection between the generality of cau-

sation and the methods used to test causal relations (see Woodward 2004). Bogen, we have seen, disagrees. The merits of and relations among these approaches have been discussed at some length by Craver (2007), Glennan (2002, 2010), Psillos (2004), Tabery (2004) and Woodward (2002, 2010). Leuridan again does not address this discussion.

One last point deserves mention before almost leaving Leuridan's putatively ontological discussion. Leuridan distinguishes between P-laws and (c)P-laws. This distinction is required for Leuridan to distinguish P-laws that are merely useful for prediction from those that, in addition, allow one to explain and control events. One might predict that one is about to run out of gas by looking at one's gas gauge, but the reading on the gas gauge does not explain the emptiness of the tank. Nor could one make it further down the road by breaking the gauge. For this reason, Leuridan (like the mechanists Glennan (2002) and Craver (2007)) appeals to Woodward's systematic theory of causation (2003). According to that theory, very roughly, causal regularities are stable regularities that continue to hold when one intervenes to change the cause variable. This

view of (c)P-regularities, however, depends fundamentally on the idea of an intervention. It also depends on the notion of an ideal intervention, which is one that intervenes *via some causal paths and not others*. It also depends on a thesis of modularity: that it is possible to intervene independently on the different components of a mechanism. As Woodward acknowledges time and again, this view of the semantics of causal claims is not intended as a reductive, metaphysical analysis of the notion of cause. It would be circular as such because one requires an antecedent notion of causation to ground these features of the account (interventions, uncontrolled paths, and modularity). Ironically, a singular notion of causation such as Bogen defends might be just what Woodward's account of intervention and modularity need for their metaphysical ground. If so, then the claim that (c)P-laws are metaphysically more fundamental than singular causation would have the story exactly backwards. But these are complicated matters that we must leave for now.

3.3 *Laws without mechanisms?*

Above we focus on Leuridan's claim that cs-mechanisms ontologically depend on macrolevel (a) as well as on microlevel (b) (c)P-regularities (L1). For the sake of completeness, let us consider Leuridan's second ontological thesis (L2), that there can be (c)P-regularities without underlying mechanisms. Leuridan needs this second thesis to establish the desired "ontological asymmetry between P-regularities and cs-mechanisms" (2010, 331). In his hands, this amounts to the claim that it is possible that there are fundamental (c)P-laws, that is, (c)P-laws for which no mechanisms exist.

Leuridan does not argue for this thesis, but it seems to us at least conceivable that the world is structured with fundamental (c)P-laws (Glennan (1996, 2002, 2010) embraces this view). To decide whether this conceivable ontological picture is actual, however, would require further argument. It is also conceivable that the world has an infinite series of mechanisms within mechanisms, or that it grounds out ultimately in individual singular causal relations (as Bogen recommends), or perhaps that it grounds out in occurrent matters of fact. Leuridan has no argument to convince us that we are in one of

these worlds rather than the other, and we therefore see no compelling reason for a mechanist to take sides.

4. Mechanism and Epistemology

Let us turn finally to Leuridan's claim that p-law statements are epistemically fundamental to mechanistic models. First, he argues that explanatory mechanistic models must include p-law statements (L3), and so mechanistic explanation cannot proceed in the absence of p-law statements.⁸ Second, he claims that mechanistic knowledge is dispensable in our search for p-laws. For instance, by

⁸ One might have expected Leuridan to defend the epistemic claim that one cannot learn about mechanisms in the absence of p-laws. One might hold that one can test causal connections only on the basis of regularities. Such a claim would be false, of course, as we might make causal inferences on the basis of temporal succession or spatio-temporal contiguity, for example. Leuridan might claim (correctly) that such inferences are fallible, but all inductive inferences are fallible, including those involving p-laws.

using randomized experimental designs, one can control for disturbing mechanistic factors without knowing what they are. This latter argument is supposed to show that although (explanatory) mechanistic models are “epistemologically dependent” on p-law statements (in the sense that the former require the latter), p-law statements are not “epistemologically dependent” on mechanistic knowledge (L4) (some p-law statements can be discovered without relying on knowledge about mechanisms). Finally, he argues that if our knowledge of laws did depend upon knowledge of mechanisms, then we would face an “infinite (and vicious) epistemological regress” (2010, 333). Because knowledge of mechanisms requires knowledge of laws, our knowledge of laws and mechanisms would never ground out in fundamental facts. Fortunately, Leuridan claims, we can know the p-laws without knowing anything about mechanisms, and this blocks the regress.

Leuridan’s first point about explanatory models derives from the above discussion (see Section 3). If cs-mechanisms require (c)P-laws, then an adequate model of the cs-mechanism requires (c)P-laws. Above, we reject the antecedent. Given that not all mecha-

nisms produce a behavior in a regular way (granted, many do) there exist cases of one-off mechanisms (Salmon/Railton mechanisms) in which the mechanistic model for the irregular behavior necessarily involves neither a macro p-law statement nor a micro p-law statement. Similarly, in cases where a mechanism behaves regularly even though this macro regularity is sustained by micro-irregularities, a mechanistic model might involve a macro p-law statement but not a micro p-law statement.

But what shall we do about the cases in which biologists explain a general phenomenon in terms of general facts about components and their activities? Mechanists should not, and do not, deny the existence of such explanations. Instead, mechanists deny that an explanatory model *must* be formulated in terms of generalizations.⁹

⁹ Thus, we deviate from the central thesis Holly Andersen (2011) defends in her short response to Leuridan's paper: "The existence of stable regularities in nature is necessary for either model of explanation: regularities are what laws describe and what mechanisms explain." (2011: 325).

One *would* think that explanatory models of cs-mechanisms must include p-law statements if one embraced a covering law (CL) model of explanation, according to which explanations are arguments that subsume descriptions of events under general law statements. No mechanist, however, accepts the CL model of explanation (see especially Salmon 1984 and Craver 2007). The reasons are too widely known to be repeated here; and it would be uncharitable to saddle Leuridan with this much-maligned view when he has not explicitly endorsed a view on the matter.

For mechanists, in contrast, mechanistic explanatory models have explanatory value in virtue of the fact that they represent the relevant portion of the causal structure of the world, not in virtue of the fact that they have a canonical representational form. Explanatory models of mechanisms might be diagrams, equations, exemplars, prototypes, texts, videos, or what have you. In each case, the representational object of the model might be singular or general. What makes these models explanatory, to the extent that they are, is that they correctly (or approximately) describe the causal structures that produce, underlie, or maintain the explanandum phenomenon. Vide-

os and diagrammatic models, for example, need not include *any* linguistic expressions or equations (cf. Perini 2005). Exemplar models, such as a labeled diagram of a single mechanism working from beginning to end, by their very nature describe representative instances rather than general types. The intended scope of the explanatory model, in other words, need not be represented in the model itself. Once one thoroughly abandons the CL model, there is no justification for demanding that explanatory models must include p-law statements. Leuridan's claim that all mechanistic models must include macro p-law statements, that is, generalizations about the behavior of the mechanism, is an *unnecessary* restriction on mechanistic models that ignores the plain fact that mechanistic models are frequently developed *without* asserting within the model that it can be generalized to other phenomena. This narrow focus blinds one (as it blinded earlier generations in the philosophy of biology) to the diversity of representational forms one finds in science. Mechanists such as Salmon (1984) and Craver (2007) have therefore rightly separated the question of explanation from the question of how explanatory knowledge is represented.

Let us now consider Leuridan's second epistemic thesis. For the record, we know of no mechanist who insists that one can test p-law statements *only if* one relies on prior mechanistic knowledge. However, let's think through Leuridan's argument. His sole example is a randomized clinical drug trial, in which subjects are randomly sorted into two groups, one of which is given a drug, and one of which is given a placebo. Leuridan perhaps should have acknowledged that the effectiveness of one's randomization procedure might depend upon the mechanism of randomization. What counts as random with respect to one experimental situation will not count as random with respect to another. Were one to survey the extent of homelessness in a geographic region by conducting a poll randomized by street address, the mechanism of randomization would be systematically biased to target people who have homes. Were one to randomize drug trials by zip code, environmental factors could confound the results. The procedure would not be random in the relevant respect. What matters is whether the apparent randomization procedure is likely to sort participants into two groups that have the same distribution of potentially confounding causal factors. Whether

the randomization procedure achieves that, depends on assumptions about the relevant causal mechanisms at play (even if those assumptions are often so obvious as to be not worth mentioning). Likewise, Leuridan might have acknowledged that the standard procedure of giving placebos to control groups reflects prior knowledge of the mechanism of the placebo effect. Indeed, experimenters generally take extreme measures to match the experimental and control groups in every way that might possibly confound the results. They test for missing control conditions by asking whether there is some possible difference between the two groups that could plausibly account for the observed changes in the result/effect. Background knowledge about possible mechanisms is often central to that task.

Consider in a bit more detail about how such experiments work. A standard experiment for testing a (c)P-law involves intervening into a putative cause variable, C, and detecting from the putative effect variable, E. Mechanistic details are often crucial for assessing the appropriateness of one's interventions. As discussed briefly above, one wants to ensure that one's intervention produces the effect in E (if any) via C and not via some other mechanism.

That is, the intervention should change C. It should not change E directly. It should not change directly the value of any variable between C and E. Furthermore, the intervention on C itself should not be correlated with any other variable that is a cause of E (unless it is causally intermediate between C and E). In some cases, one wants to ensure that the intervention severs the causal influences of other variables on C so that one can attribute any change in E to the intervention alone. All of these assumptions behind the use of interventions to test (c)P-laws are assumptions about the causal structure, the mechanisms, involved in the intervention technique and in the system under study. An adequate philosophy of experimental intervention thus might make considerable progress by asking how mechanistic knowledge enters into these test procedures (see Woodward 2003; see summary diagram in Craver 2007, Ch. 3).

What about the detection component of a test for a (c)P-law? Allan Franklin (2009) has generated a useful list of strategies by which scientists confirm that their techniques are reliable indicators of phenomena such as E. Many of these strategies rely crucially on facts about the mechanisms at play. One might, for example, argue

that there could be no other cause of the measured value of E besides the fact that E has that value. One might show that one's technique reliably registers reliable artifacts known to be produced under aberrant causal conditions. One might rely on a theoretical understanding of the mechanism by which the detection technique works. One might check the results of one's technique against another technique that relies on causally independent mechanisms (see Franklin 2009). In each of these cases, one relies on knowledge about the mechanisms involved in the system and in the detection technique to argue that the methods in question provide an adequate measure of E in these circumstances. In short, even if it is possible to test (c)P-laws without knowing mechanisms (and we deny that Leuridan's example shows as much), one might learn a great deal about how (c)P-laws are tested by thinking about the mechanisms involved in the test conditions. By casting the debate as a forced choice between laws and mechanisms, one occludes far more interesting questions about how mechanistic knowledge contributes to the design and interpretation of experiments for testing p-laws.

Finally, Leuridan claims that if our ability to test (c)P-laws relies exclusively on cs-mechanisms, then we face an infinite regress. The regress arises because if cs-mechanistic knowledge relies upon knowledge of (c)P-laws (or, more precisely, mechanistic explanations must involve p-law statements), and if knowledge of (c)P-laws requires knowledge of cs-mechanisms, then we never reach the epistemic bottom. It is not clear to us that this is a well-formed problem, and so we are not clear how to solve it.¹⁰ We know of no foundationalist who proposes to build scientific knowledge out of the basic building blocks of mechanisms or laws. Foundationalists tend to construe the epistemic foundation of science in terms of particular matters of fact, sense data, or innate ideas, not in terms of p-laws or mechanisms. We think that both p-laws and mechanisms contribute to the advancement of science, and we

¹⁰ Contrary to Leuridan's claim, Machamer, Darden, and Craver (2000) discuss bottom-out activities not as a way of solving some sort of epistemic regress but as a disciplinarily relative way of identifying when explanations come to an end.

feel no pressing need (and have been given no compelling argument) to place one above or below the other in the order of our knowledge. Furthermore, it is not at all clear from Leuridan's formulation how laws stop the regress. If one must know p-laws in order to adequately test p-laws (for example, p-law statements that one's randomization procedure regularly randomizes, that one's interventions work the same way each time, and so on), then one still has a regress of sorts, and Leuridan has not shown how it will come to an end. How can we design a randomized experiment if we cannot trust that our randomizing procedure generally randomizes? And how can we control for confounding factors if there are no general facts about which factors are confounding? How do we know that our intervention is adequate if there are no general facts about how our intervention works? It would appear that laws are no more epistemically secure than are mechanisms in the foundationalist view that Leuridan apparently embraces.

5. Conclusion

For the discussion of these matters to move forward, it is crucial not to manufacture an artificial conflict between philosophers who emphasize the centrality of mechanisms in our thinking about science and philosophers (such as Mitchell) who seek a plausible way to talk about generalization in science. No mechanist denies that there are pragmatically useful regularities. And nobody who thinks there are pragmatically useful regularities should feel any pressure to deny that the search for mechanisms is central to the practice of biology and many other sciences.

It is a surprising fact about the history of the philosophy of science that of these two correlative concepts, generalizations have tended to dominate the discussion. Against this backdrop, mechanists should be read as suggesting something of a gestalt-shift in which mechanisms are moved into the foreground. Such a shift leads attention away from the formal structure of scientific theories (and questions about the logical structure of law-statements and models) and toward the material structures that scientists endeavor to describe. Attention to such material structures provides resources for

thinking about how generalizations and mechanisms are discovered, evaluated, and extrapolated and into how such concepts are deployed in explanation, prediction, and control. The perceived need to defend laws, no matter how much they have been weakened and stripped of their once-robust metaphysical content, reflects a conservative refusal to acknowledge that perhaps the philosophy of science might benefit from coming at its subject matter from a fresh perspective. Mechanists decenter laws in their thinking about science because the old paradigm, centering laws, has become mired in debates that are inconsequential and, as a result, have stopped generating new questions and producing new results. In this paper, we have argued that by trying on the mechanistic gestalt, one can make progress on problems concerning explanation, laws, prediction, and manipulation where the nomic approach seems to have run out of gas. Moving forward, there are far more interesting and better-motivated questions to ask than whether mechanisms can replace generalizations or vice versa.

References

- Allen, Garland E. 1978. *Thomas Hunt Morgan: The Man and His Science*. Princeton: Princeton University Press.
- Andersen, Holly K. 2011. Mechanisms, Laws, and Regularities. *Philosophy of Science* 78: 325-331.
- Beatty, John. 1995. The Evolutionary Contingency Thesis. In *Concepts, Theories, and Rationality in the Biological Sciences*, ed. Gereon Wolters, and James G. Lennox, 45-81. Pittsburgh: University of Pittsburgh Press.
- Bechtel, William. 2009. Generalization and Discovery by Assuming Conserved Mechanisms: Cross Species Research on Circadian Oscillators. *Philosophy of Science* 76: 762-773.
- Bechtel, William, and Adele Abrahamsen. 2005. Explanation: A Mechanist Alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 421-441.
- Bechtel, William, and Robert C. Richardson. 2010. *Discovering Complexity. Decomposition and Localization as Strategies in Scientific Research*. Cambridge: MIT Press.

- Bogen, James. 2005. Regularities and Causality; Generalizations and Causal Explanations. *Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 397-420.
- — —. 2008. Causally Productive Activities. *Studies in History and Philosophy of Science* 39: 112-123.
- Craver, Carl F. 2007. *Explaining the Brain. Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Clarendon Press.
- — —. (2002) Structures of Scientific Theories. In *Blackwell Guide to the Philosophy of Science*, eds. P.K. Machamer and M. Silberstein, 55-79. Oxford: Blackwell.
- Darden, Lindley. 1991. *Theory Change in Science: Strategies from Mendelian Genetics*. New York: Oxford University Press.
- — —. 2006. *Reasoning in Biological Discoveries: Essays on Mechanisms, Interfield Relations, and Anomaly Resolution*. New York: Cambridge University Press.
- Darden, Lindley, and Carl F. Craver. 2002. Strategies in the Interfield Discovery of the Mechanism of Protein Synthesis. *Studies in History and Philosophy of Biological and Biomedical Sciences* 33: 1-28.

- Franklin, Allan. 2009. Experiment in Physics. *The Stanford Encyclopedia of Philosophy (Spring 2010 Edition)*, ed. Edward N. Zalta, URL = <http://plato.stanford.edu/archives/spr2010/entries/physics-experiment/>.
- Goodman, Nelson. 1955. *Fact, Fiction, and Forecast*. Cambridge: Harvard University Press.
- Glennan, Stuart S. 1996. Mechanism and the Nature of Causation. *Erkenntnis* 44: 49-71.
- — —. 2002. Rethinking Mechanistic Explanation. *Philosophy of Science* 69: 342-353.
- — —. 2010. Mechanisms, Causes, and the Layered Model of the World. *Philosophy and Phenomenological Research* 81: 362-381.
- Hull, David L. 1978. A Matter of Individuality. *Philosophy of Science* 45: 335-360.
- Leuridan, Bert. 2010. Can Mechanisms Really Replace Laws of Nature? *Philosophy of Science* 77: 317-340.

- Machamer, Peter. 2004. Activities and Causation: The Metaphysics and Epistemology of Mechanisms. *International Studies in the Philosophy of Science* 18: 27-39.
- Machamer, Peter, Lindley Darden, and Carl F. Craver. 2000. Thinking About Mechanisms. *Philosophy of Science* 67: 1-25.
- Mitchell, Sandra D. 1997. Pragmatic Laws. *Philosophy of Science* 64 (Proceedings): 468-479.
- — —. 2000. Dimensions of Scientific Law. *Philosophy of Science* 67: 242-265.
- — —. 2003. *Biological Complexity and Integrative Pluralism*. Cambridge: Cambridge University Press.
- — —. 2009. *Unsimple Truths. Science, Complexity, and Policy*. Chicago/London: University of Chicago Press.
- Morgan, Thomas H., Alfred Sturtevant, Hermann Muller, and Calvin Bridges. 1915. *The Mechanism of Mendelian Heredity*. New York: Henry Holt and Company.
- Perini, Laura. 2005. Explanation in Two Dimensions: Diagrams and Biological Explanation. *Biology and Philosophy* 20: 257-269.

- Psillos, Stathis. 2004. A Glimpse of the *Secret Connection*: Harmonizing Mechanisms with Counterfactuals. *Perspectives on Science* 12: 288-319.
- Salmon, Wesley. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
- Schaffer, Jonathan. 2000. Causation by Disconnection. *Philosophy of Science* 67: 285-300.
- Schaffner, Kenneth F. 1993. *Discovery and Explanation in Biology and Medicine*. Chicago/London: University of Chicago Press.
- — —. 2008. Theories, Models, and Equations in Biology: The Heuristic Search for Emergent Simplifications in Neurobiology. *Philosophy of Science* 75: 1008-1021.
- Smart, J.J.C. 1963. *Philosophy and Scientific Realism*. London: Routledge.
- Steel, Daniel P. 2008. *Across the Boundaries. Extrapolation in Biology and Social Science*. Oxford: Oxford University Press.
- Tabery, James G. 2004. Synthesizing Activities and Interactions in the Concept of a Mechanism. *Philosophy of Science* 71,1-15.

- Wimsatt, William C. 1998. "Simple Systems and Phylogenetic Diversity." *Philosophy of Science* 65: 267-275.
- Woodward, James. 2002. What Is a Mechanism? A Counterfactual Account. *Philosophy of Science* 69: 366-377.
- — —. 2003. *Making Things Happen. A Theory of Causal Explanation*. Oxford: Oxford University Press.
- — —. 2004. Counterfactuals and Causal Explanation. *International Studies in the Philosophy of Science* 18: 41-72.
- — —. 2010. Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation. *Biology and Philosophy* 25: 287-318.