

A field guide to mechanisms: part I

Holly Andersen
Simon Fraser University
handerse@sfu.ca

Abstract: In this field guide, I distinguish five separate senses with which the term ‘mechanism’ is used in contemporary philosophy of science. Many of these senses have overlapping areas of application but involve distinct philosophical claims and characterize the target mechanisms in relevantly different ways. This field guide will clarify the key features of each sense and introduce some main debates, distinguishing those that transpire within a given sense from those that are best understood as concerning distinct senses. The ‘new mechanisms’ sense is at the center of most of these contemporary debates, and will be treated at greater length; subsequent senses of mechanism will be primarily distinguished from this one. In part I of this paper, I distinguish two senses of the term ‘mechanism’, both of which are explicitly hierarchical and nested in character, such that any given mechanism is comprised of smaller sub-mechanisms, in turn comprised of yet smaller sub-sub-mechanisms, and so on. While both of the senses discussed here are anti-reductive, they differ in their focus on scientific practice versus metaphysics, in the degree of regularity they attribute to mechanisms, and in terms of their relationships to the discussions of mechanisms in the history of philosophy and science.

Keywords: mechanisms; causation; explanation; reduction; methodology

0) Introduction

Talk of mechanisms is central to a variety of recent debates in philosophy of science. For those working outside of these debates, it can be difficult to get a clear sense of what all the talk of ‘new mechanisms’ is about. This two-part field guide is intended to provide an overview of the main new sense of mechanism that has been the primary focus of discussion, and to distinguish this new sense from other related but relevantly different senses in which the term mechanism appears in philosophy of science. This field guide begins by a concise overview of the literature on ‘new mechanisms’, and then turns to distinguishing and comparing other notions of mechanism to this one.

This field guide offers taxonomic classifications of species of mechanism, with key distinguishing markings of each species. These taxonomic distinctions have not been utilized in debate involving 'mechanisms' but are motivated out of the kinds of characterizations and debates that have multiplied on this topic recently.

In elaborating each distinct sense, I am thereby arguing that each involves a unique set of ontological, methodological, explanatory, and anti-/reductive commitments. Claims that are true of one sense will fail of others; as such, the senses are not interchangeable, even though they are often presumed to be (especially mechanism₁ and mechanism₂).

The division between parts I and II classifies types of mechanisms in terms of whether or not they are explicitly anti-reductive. Here in part I, I focus on two senses that include an explicitly tiered ontology as an alternative to reduction. These approaches to mechanisms concern ways to *individuate* and *connect* levels, especially but not limited to levels as they are studied in the sciences, while rejecting that this constitutes a *reduction* of higher levels to lower levels. This anti-reductive approach to mechanisms aims to provide explanatory autonomy to the so-called special sciences, while still grounding their objects of investigation firmly in the physical world. These two senses can be distinguished by differing in 3 major characteristics: the first emphasizes the role of mechanisms in various aspects of scientific practice, while the second emphasizes a metaphysical role for mechanisms; the first attributes some degree of regularity to mechanisms, while the second allows any unique causal chain to count as a mechanism; and while the first offers something novel to the ongoing philosophical discussion of mechanisms, the second is more closely aligned with the historical notion of mechanism. These characteristics are mutually supportive; for instance, an emphasis on scientific practice leads clearly to a focus on regularities as targets for experimentation, while an emphasis on metaphysics connects much more closely to the historical discussion of mechanisms in natural philosophy.

Each section of this field guide will begin by briefly characterizing the most salient philosophical features of the relevant species of mechanism, including the kind of discussion in which it most frequently appears. I will then briefly discuss the ontological commitments of each sense, their methodological implications, and their role in explanation, respectively. Each section includes a comparison distinguishing it from other related senses of mechanism.

1) Mechanism₁: Mechanisms as integral to scientific practice

This is the ‘new mechanism’ approach that is arguably the primary sense of the term mechanism in recent discussion. This species of mechanism is very commonly found in discussions of explanation in the sciences, as distinct, for instance, from explanation in history or of action. It comprises several claims. One is the ontological claim that the phenomena studied by many sciences, especially so-called higher-level sciences like biology, have a particular hierarchical structure comprised of nested levels. Each level involves entities organized in various ways and connected via causal interactions or activities, constituting or giving rise to a pattern that recurs under specific identifiable circumstances. These patterns are the regularities that constitute the phenomena studied by these sciences. Such mechanisms provide non-arbitrary ways to divide parts of the world into levels of physical size and organization, and to understand the connections between different levels, without having to reduce or eliminate higher levels or treat them as merely apparent or epiphenomenal.

A second claim is that the methodology of many sciences reflects this ontological structure of their target phenomena, by isolating consistent patterns or regularities in the world and then decomposing them into constituent entities, causal connections, and spatio-temporal organization. The emphasis is on the practice of science. A third claim, central to this sense, concerns explanation in such sciences,

namely, that the explanations they offer reflect this ontology. Explanations are of regularly recurring phenomena in the physical world that are the end product of, or are constituted by, the operation of such mechanisms; such regularities are to be explained by providing details about the mechanism(s) responsible for producing them.

The trajectory leading up to this sense of mechanism is familiar. As logical empiricism waned in the middle of the 20th century, philosophers of science focused on deduction from universal laws as the primary form of explanation in science (e.g. Hempel and Oppenheim 1948, Hempel 1963). On this view, an explanation involves a general law plus a set of initial conditions, from which the explanandum could be deduced.¹ A variety of problems plagued the deductive-nomological account of explanation (see, for instance, Salmon 1978). One clear issue was the apparent lack of suitable laws in biology to serve in a deduction. This posed an unfortunate dilemma: either biology and related sciences do not really explain, or the D-N account applies only narrowly to some explanations, mainly in physics. Laws can be reconstrued in ways that render them more applicable in biological sciences (Mitchell 1997), but many still found the D-N model of explanation unsatisfactory in biological contexts.

This 'new mechanism' approach arose in response to many of the issues raised with the D-N model of explanation. It differed from the D-N account in terms of what it took the sciences to be investigating, how those investigations generally proceed, and the explanatory results of those investigations. Ontologically, it shifted the emphasis away from universal laws and nomic necessity, towards recurrent patterns of causal structure, with a much more limited scope of generalizability. Methodologically, it took the sciences to be engaged in finding more and more detail, by targeting and decomposing the mechanisms that support limited regularities,

¹ On this view, the deductive arguments that rely on universal laws to show the necessity of the explanandum also have an inductive counterpart, relying on statistical generalizations that raise the probability of the explanandum.

rather than seeking ever broader and more abstract generalizations. These ontological and methodological views support an account of explanation that eschews propositionally structured laws and logically valid arguments, instead relying on mechanism models, of varying levels of descriptive specificity, involving real entities, their activities, and their organization into coherent processes that recur under appropriate circumstances.

Finally, this new sense of mechanism is deeply anti-reductionist: science may uncover explanatory or ontological *connections* between higher and lower levels, but does not thereby either eliminate or *reduce* the higher levels thus connected (Wimsatt 1994). The hierarchically structured, anti-reductionist approach involved in mechanism₁ holds that science provides genuine explanations that “bottom out” without having to be grounded in some ultimate fundamental level (Machamer, Darden, and Craver 2000). Higher-level explanations are not viewed as placeholders for some ‘real’, microphysical, explanation; the focus is instead on integrating and connecting levels (Glennan 1995, 2002, 2005; Darden and Craver 2009). Even if mechanisms in biology are ultimately composed of sub-sub-submechanisms in physics, this does not mean that biological phenomena explained by such mechanisms are thereby reduced to the lower-level mechanisms.

This anti-reductive stance has ontological consequences. Methodologies for investigating levels of recurrent mechanisms can proceed in relative autonomy, neither affirming nor denying any ultimate ontological relationship between different branches of science, scientific theories, or phenomena. Mechanism₁ thus imposes some distance between, on the one hand, normative epistemological or methodological claims about the most effective ways to investigate the world, and, on the other hand, the underlying nature of that world. It is not explicitly anti-metaphysical, but rather metaphysically agnostic. The anti-reductive character of mechanism₁ allows us to make methodological recommendations about investigating the world (see especially Bechtel and Richardson 1992 and Darden

2002) without thereby committing ourselves to a single account of what that world is like.

Mechanism₁ is often elaborated as part of a detailed look at the actual investigatory practices of the sciences, especially biological sciences. Bechtel and Richardson (1993; see also Bechtel and Abrahamsen 2005) characterize scientific discovery in the biological sciences in terms of localizing and decomposing mechanisms. Instead of searching for universal laws, many scientists are engaged in the task of identifying regular phenomena and the precise circumstances under which they transpire, and the locus of control by which mechanisms can be manipulated. Some target regularity must first be delineated – what is the phenomenon to be explained? This often proceeds by investigating the precise spatio-temporal boundaries of the regularity, including parameter values that modulate how it occurs and supporting environmental factors outside of the mechanism itself. Localized mechanisms can be decomposed into constituent elements or components, their interactions, and the organization of those elements into a coherent and re-identifiable process. Once localized and decomposed, the delineated mechanism explains why the identified regularity occurs as it does, under the conditions that it does.

The way in which mechanisms₁ are investigated shapes the kinds of explanations in which those mechanisms figure. Explanations can involve a greater or lesser degree of specificity, depending on the state of scientific knowledge or the context in which a given explanation is provided. For example, Craver and Darden (2001), Darden (2002), Craver (2007), and Darden (2008) offer mechanism sketches, mechanism schemata, and fully developed mechanisms as identifiable stages involved in the discovery of mechanisms.

Mechanism schemata are abstract descriptions of mechanisms that can be instantiated to yield descriptions of particular mechanisms. The term *mechanism schemata* is fitting because their components are placeholders that

can be filled in with detailed stages between the setup and termination. ... Mechanism sketches, in contrast to mechanism schemata, are abstract descriptions of mechanisms that cannot yet be filled in (Craver and Darden 2001, 120).

Darden calls these distinctions “advisory”: neither simply descriptive, nor purely prescriptive, they provide a set of terms and concepts for scientists that can be genuinely useful in structuring research, while recognizing that these terms and concepts have been developed by considering what many scientists are already doing.

Schemata and sketches vary in terms of the amount of detail they fill in about the mechanism(s) in questions and the amount of possibility they leave open for how that mechanism might be instantiated. The idea is that a mechanism sketch places constraints on the kind of mechanism that could instantiate it. The first stage of discovery of a mechanism may consist in such a sketch that guides further research by narrowing the range of possible kinds of mechanisms to be investigated. A fuller set of constraints may be provided by mechanism schemata, where components and their connecting activities have sufficient constraints placed on them that researchers can begin filling in physical details for particular mechanisms. A single schema may be instantiated by several distinct mechanisms that fill in the placeholder components with different entities or activities that fulfill broadly similar functions.

This highlights the need to distinguish between mechanisms₁, which are in the world, and models of mechanisms₁ (e.g. Glennan 2005, Illari and Williamson 2010). This distinction figures in the debate of what it is in virtue of which mechanisms₁ are explanatory. Drawing from Salmon (1984), the ontic conception of mechanisms holds that mechanisms are *represented* in models and theories in the sciences, but are not identical with those representational devices. They are instead ‘out there’ in the world. This makes an important difference when considered in terms of

explanation: do semantic or syntactic relationships between elements of a model-theoretic representation do the explanatory work in their own right, or do they merely point to or exhibit the causal relationships in the world that do the explanatory work? Wright (2012) argues against this ontic conception, holding instead that explanation is a communicative process for the purpose of understanding which must thus involve models of mechanisms. Many models of mechanisms₁ are not propositionally structured, but are instead conveyed using diagrammatic or visual forms of representation (Darden and Craver 2002, Perini 2005a, 2005b). This debate concerns the role that mechanism₁ plays in explanation: do mechanism explanations involve physical mechanisms themselves, or models of those mechanisms? It has consequences that reach beyond the discussion simply of mechanisms₁, however, by undermining general accounts of explanation that rely on, for instance, logical entailment relationships between premises in arguments to characterize explanation (Strevens 2008). If mechanism₁ models figure in explanations, diagrammatic or other visual representations of mechanisms₁ lack the requisite propositional structure to figure in such logical relationships; if physical mechanisms₁ themselves figure, they lack such features even more obviously.

The extent to which mechanisms₁ constrain ontological commitments is also a point of dispute. Some of these debates are best understood in terms of a univocal sense of mechanism (namely, mechanism₁), while others involve tension between different senses of mechanism. Two examples of intra-sense debate concern how to best characterize the internal structure of mechanisms, and the extent of ontological commitments entailed by mechanism₁. For the first example, there are disagreements about how to best describe the process of a mechanism in action. Machamer, Darden, and Craver (2000) utilize start-up conditions that trigger the mechanism, and termination conditions that are the last stage within the mechanism. Using the example of circadian rhythms, Bechtel and Abrahamsen (2010) argue instead that mechanisms are often cyclical, and therefore lack clear start-up or termination conditions.

The second example involves disagreement about how much ontological commitment should accompany mechanism₁. While some authors who defend this new mechanism approach do so in an ontologically agnostic way (for instance, Bechtel and Richardson 1993), mechanism₁ is clearly well-suited to certain kinds of ontological commitments.² For instance, Machamer, Darden, and Craver (2000) argue for a dualistic entity-activity ontology to ground the mechanistic₁ account of explanation. They characterize mechanisms as “entities and activities organized such that they are productive of regular changes from start or set-up conditions to finish or termination conditions” (2000, 3). Mechanisms₁ are organized chains of entities connected via productive activities of those entities; the organized activities provide both continuity to the mechanism and regularity of operation (ibid.). Phenomena to be explained via mechanisms are either the final product that results once termination conditions are reached or a higher-level regularity to which the mechanism gives rise. Thus, mechanisms can explain (at least) two ontological kinds of target phenomenon. One kind is the effects of causal relationships, where the explanandum is the final stage in an often complex causal chain at the same level as the rest of the entities and activities in the mechanism. The second kind is phenomena at a higher level than the explanans, where causal activities connect the entities at a lower level, and the entire mechanism’s operation constitutes the higher-level phenomenon.

Machamer, Darden, and Craver (2000) explicitly aim to offer a view of mechanisms₁ that is compatible with Salmon’s (1984) process account of causation. In contrast to their commitment to the process account, the causal relationships binding mechanisms₁ together could be construed in terms of counterfactuals, or dispositions, or capacities (e.g. Ilari and Williamson 2010). Accordingly, we see that mechanism₁ is not itself an account of causation (this will be more fully explored in

² This is not to say that mechanisms₁ are compatible with just any ontology: the epistemological or methodological features of the new mechanisms approach place compatibility constraints on ontological commitments. Rather, the point here is that these constraints are not enough to dictate a single unique ontology.

the next section), however much claims about causation figures in ongoing debates within the mechanisms literature. Williamson's (2011) criterion for being a mechanistic theory of causality, that "two events are causally connected if and only if they are connected by an underlying physical mechanism of the appropriate sort" (421), is not met by mechanisms₁.

The anti-reductive aspect of mechanisms₁ generates questions about interlevel causation; this is a debate concerning causal relationships within mechanisms₁, not about the nature of causation per se. The nested, hierarchical structure of mechanisms allows for non-arbitrary ways to individuate levels in the world, in a local and context-sensitive rather than broad or generally applicable way. It also raises the question of whether direct causal relationships between entities or processes at different levels are possible.³ Craver and Bechtel (2007) reject the possibility of interlevel causation. They argue for a differentiation of levels such that it is only meaningful to speak of higher and lower levels within a single given mechanism. The levels in a mechanism form a hierarchy structured by the constitutive relationships between phenomena and the mechanisms that give rise to them. Craver and Bechtel note a common stricture on causal relata, such that parts of a whole cannot directly enter into causal relationships with that whole. They conclude that there is no meaningful way to speak of causal relationships between different levels within a mechanism. Either those relationships must be non-causal and constitutive, or they must not be between genuinely different levels, since they occur between different mechanisms and levels are only defined within individual mechanisms.

This conclusion, however, is in tension with much scientific practice, where levels are often treated as robust differentiations across multiple mechanisms. Mitchell (2008) provides a number of examples of what she calls multilevel causation and

³ There are many ways to construe levels; Craver (2007) offers a different kind of field guide to many of those ways. Here, I focus on levels within a mechanism, which are defined in terms of nested mechanisms.

explanation, where causal relationships connect levels in complex systems. This complexity of causal interaction is reflected in the actual explanatory practices of the sciences, as she illustrates with case studies. This puts the onus on those who would deny interlevel causal relationships to explain away the apparently widespread attribution of such relationships by scientists who work on complex systems.

The dispute about interlevel causation and mechanisms₁ can be resolved by clarifying a wide-spread type-token ambiguity with respect to mechanisms₁. An apparently integral part of why mechanisms₁ are so useful in the sciences is that an explanation involving a mechanism₁ might involve a token of a mechanism, or of a type of mechanism; it might explain a token occurrence of a phenomenon, or explain a type of phenomenon. Explanations providing a mechanism are often ambiguous regarding this distinction. A mechanism₁ may indicate a type of mechanism, such as a calcium ion channel, or a particular mechanism, such as *this* calcium ion channel. Explanations of token instances of mechanisms₁ explain those occurrences by referring to them as of a particular type, namely, the general mechanism₁ type. Sometimes the target explanandum is just one single channel in one cell in vitro in a lab. But more often, it is the regularly recurring phenomenon that is displayed in multiple instances of channels, in any such channel in any cell. The generality of the mechanism does the explanatory work in the individual cases, while general explanations are developed via individual cases. This allows us to disentangle two different claims regarding interlevel causation. Types of sub-mechanisms at different levels within a single nesting mechanism-type may violate the stricture against self-causation, while different tokens of those same sub-mechanism types might be perfectly acceptable causal relata.

This brings us to a key feature of mechanisms₁ by which they can be distinguished from other senses, especially from mechanisms₂. Mechanisms₁ are characterized in part by *regularity*. Mechanisms₁ are the kinds of things that do or straightforwardly could recur. Accounts of mechanisms₁ are developed with an eye towards contemporary scientific practice, in what are often called the special sciences,

focused on methodology, experimentation, and explanatory practices in those sciences. In this regard, mechanisms are agnostic with respect to, for instance, physics. It may be that the deductive-nomological model best accounts for explanatory practices in fundamental physics; the criticisms of the D-N model that led to the new mechanisms approach highlight how that model is inadequate for *some* sciences, in terms of an alternative that works best within those sciences. The regularity involved in mechanism₁ ties it to the sciences. Rather than studying unique, singular causal events, the sciences study phenomena that are repeatable, or studies unique events (such as speciation events) as potentially repeatable.

This close tie to scientific practice, however, disconnects mechanism₁ from earlier discussions of mechanism. Historically, mechanisms were part of a complete mechanistic worldview: the entire world (or at least, the material parts of it) was viewed as mechanistic in character. The new mechanisms literature offers a view of mechanism₁ that differs from this historical discussion with respect to the completeness of the associated ontology. Because mechanism₁ guide and structure research methodology and scientific explanations, they provide constraints on our ontological commitments: it should turn out that those scientific practices largely constitute a well-established epistemological method for investigating the structure of the world. But scientific practice does not constrain enough to pronounce one or another account of ontology or causation as the one to which we should be committed; it leave room for agnosticism in these regards. For a complete ontology based on mechanisms, including explanations of singular events as singular, one looks to mechanism₂ instead.

2) Mechanism₂: Mechanisms as an ontology of the world

This species of mechanism₂ has re-entered the discussion in philosophy of science as part of the discussion of mechanism₁, but has several features that render it

more than a mere variation on mechanisms₁. Mechanism₂ provides a metaphysical picture of the world as comprised of hierarchically structured layers of mechanisms, each mechanism constituted by a layer of smaller mechanisms and their entities and interactions, and partially constituting further layers. None of these layers need be fundamental, in terms of being the 'real' one from which the rest are derivative. Rather, the ontology involves taking the whole hierarchy of constitutively structured layers of causation, each of which is genuinely real, and none of which need be treated as most basic (Glennan 2010a). Elucidating the mechanism(s)₂ that instantiates a particular higher-level property is not equivalent to reducing the higher-level property to a lower-level mechanism.

The ontological commitments of mechanism₂ are well-defined: it aims to provide a complete ontology of what there is, along with a complete account of causation. Mechanism₂ fits Williamson's (2011) description of what he calls a complex-systems account of mechanistic causality. Ontologically, any causal chain in the world, at any level, counts as a mechanism, even if it could in principle only occur once. When coupled with the anti-reductive focus, mechanisms₂ also provide an account of causation. Each mechanism involves causal interactions between entities within it; when we look closer at these causal interactions, we find a sub-mechanism that gives rise to it. This yields a richly causal layered view of the world (Glennan 2010a), where it's mechanisms all the way down.

Corresponding to the stronger ontological commitments of mechanism₂ is a weaker set of methodological implications for how the world so pictured should be investigated. Unlike the advisory details offered to scientists by Darden, Craver, Bechtel, Richardson, and others in the previous section, mechanisms₂ do not yield specific guidance regarding scientific investigation, other than the broad understanding that the world has this kind of layered causal and constitutive character. Finally, mechanism₂ offers a reason why some explanations are explanatory: namely, because they involve models of mechanisms₂ that track or pick out the correct underlying structure of the world.

With respect to explanation, mechanism₂ is a generalization of mechanism₁, minus the regularity requirement such that target phenomena recur under appropriate conditions. Because any causal chain will meet the criteria to be mechanistic₂ in character, there need be nothing fundamentally different about explaining unique events that cannot be repeated, such as major historical occurrences, and explaining events in terms of types that regularly recur, such as firings of neurons. While there are many authors who espouse a mechanistic₂ worldview, one recent example is particularly illuminating for the way in which it explicitly connects mechanism₁ with mechanism₂. Glennan (2010b) expands and modifies an account of mechanism₁ in the sciences in order to provide a single account of explanation that includes both scientific explanations of regularities as well as what he calls historical explanations of singular causal events. He characterizes the difference between scientific and historical explanations purely in terms of regularity: historical explanations are for phenomena that only occur once, while scientific explanations are for phenomena that occur more than once. Thus, there could be historical explanations of singular events in science, and scientific explanations of recurrent historical events. Glennan argues that the similarity between them is based on the fact that we can still provide a series of causal events, described at the right grain, in order to explain either kind of event. Rather than simply calling these explanations causal relationships, leaving open the character of causation that is involved, we should say they involve mechanisms₂. This means that, compared with mechanism₁, mechanism₂ strengthens the ontological commitments by providing a complete account of causation, ruling out counterfactual or difference-making accounts of causation.

While this in many ways offers a smooth transition from mechanism₁ in scientific explanation to mechanism₂ in explanation generally, it crosses a divide with respect to the sense of mechanism at play. The existence of some kind of regularity requirement is the key to distinguishing mechanism₂ from mechanism₁. When some sort of regularity requirement is in place, there will exist causal relationships that are not part of any mechanism₁ (Andersen 2011). Such 'historical' cases simply fall

outside the ambit of explanation by mechanisms₁. If one removes the regularity requirement, mechanisms can be used as the basis for a complete ontology of the world and account of causation. In doing so, however, they lose much of what made them useful as a characterization of scientific practice. With few exceptions, mechanisms are explanatorily powerful in the sciences precisely because they recur: regular phenomena are the main explanatory targets, and explanations pointing to those regularities are powerful because they work for so many token occurrences. Strategies for mechanism discovery such as those offered by Bechtel, Richardson, Darden, and Craver would not work for ephemeral mechanisms₂, since they cannot be the object of repeated experiments and are only tokens, not types (see discussion in previous section). Characterizing mechanisms₁ in terms of their role in the sciences requires a different account than characterizing all causal relationships as mechanistic₂ in character.

The close connection between explanation and causation is widely recognized; thus, it is somewhat counterintuitive and philosophically informative to find that there is a meaningful distinction between mechanisms₂ as an account of causation and ontology and mechanisms₁ as an account of scientific explanation.

3) Conclusion

Mechanisms in the two senses discussed in this paper provide a layered view of phenomena that rejects the reductionism that was the hallmark of much of philosophy of science during the 20th century. Both mechanism₁ and mechanism₂ are characterized in terms of entities and the causal activities that connect them at a level, plus constitutive relationships that nest upwards and downwards. They are primarily distinguished from each other by whether or not the mechanisms are regular or general. If they are, then they are mechanisms₁, and will be further characterized by an emphasis on scientific practice, providing methodological

recommendations for investigation, playing an important role in the provision of scientific explanations, and a compatibility with ontological agnosticism. If they are not required to be regular, they are mechanisms₂, and can provide a complete ontology of the world in term of layered causal chains, with a focus on metaphysics rather than details of scientific practice, and will be suitable for use in many kinds of explanation, scientific as well as historical.

We've seen that mechanisms₁ do not meet Williamson's (2011) criteria to be a mechanistic theory of causation, while mechanisms₂ do. Williamson notes how mechanistic views of the world stretch back at least to the early modern period. This is revealing of something very important in the contemporary discussion of mechanisms in philosophy of science. Mechanisms₁ are really a novel development, as a focused usage of the term for a proper subset of causal relationships in the world. By contrast, mechanisms₂ have a much closer relationship to the historical discussion of mechanisms, and arguably constitute a contemporary continuation of the early modern tradition.

The different relationships to causation have further ramifications for how each sense figures in, for instance, explanation. In section 2, I noted an ambiguity in mechanisms₁ with respect to types versus tokens: when we offer a mechanism to explain some event, it may not be specified whether the mechanism doing the explaining is a token instance or the general type. This ambiguity is not present in mechanisms₂. This ambiguity with respect to types versus tokens leads to a range of further disagreements about the relationship between mechanisms and causation.

This taxonomic distinction between mechanisms has substantial philosophical consequences for some of the debates discussed in this paper. For instance, the debate about the role of regularity in mechanisms concerns two distinct senses of the term, while the debate about start-up and termination conditions versus cycles occurs within a single sense. From outside philosophy of science, it can be easy to assume that the 'new mechanisms' that are the focus of so much discussion bear a

much closer relationship to the historical discussion than they actually do. For that, one should look to mechanisms₂.

References

- Andersen, Holly (2011). "The case for regularity in mechanistic causal explanation." *Synthese Online First* DOI 10.1007/s11229-011-9965-x.
- Bechtel, William and Adele Abrahamsen (2005). "Explanation: A Mechanistic Alternative." *Studies in History and Philosophy of Biological and Biomedical Science* 36: 421-441.
- Bechtel, William and Adele Abrahamsen (2010). "Dynamic mechanistic explanation: computation modeling of circadian rhythms as an exemplar for cognitive science." *Studies in the History and Philosophy of Science* 41(3): 321-333.
- Bechtel, William, and Richardson, Robert (1993). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Princeton University Press.
- Bogen, Jim (2005). "Regularities and Causality; Generalizations and Causal Explanations." *Studies in History and Philosophy of Biological and Biomedical Science* 36: 397-420.
- Cartwright, Nancy (2002). "From causation to explanation and back." In *The Future for Philosophy* ed. Leiter, London: Oxford Univ. Press: 230-45.
- Craver, Carl (2007). *Explaining the Brain*. Oxford University Press.
- Craver, Carl, and Bechtel, William (2007). "Top-down causation without top-down causes." *Biology and Philosophy* 22:547-563.
- Craver, Carl and Darden, Lindley (2001) "Discovering mechanisms in neurobiology: the case of spatial memory." In *Theory and Method in the Neurosciences* eds. Machamer, Grush, and McLaughlin, University of Pittsburgh Press: 112-137.
- Darden, Lindley (2002). "Strategies for Discovering Mechanisms: Schema

- Instantiation, Modular Subassembly, Forward/Backward Chaining." *Philosophy of Science* 69(S3): S354-S365.
- Darden, Lindley (2008). "Thinking Again about Biological Mechanisms." *Philosophy of Science* 75(5): 958-969.
- Darden, Lindley, and Craver, Carl (2009). "Reductionism in Biology." Wiley eLS Online Library; DOI: 10.1002/9780470015902.a0003356.
- Glennan, Stuart S. (1996). "Mechanisms and The Nature of Causation." *Erkenntnis* 44: 49-71.
- Glennan, Stuart S. (2002). "Rethinking Mechanistic Explanation." *Philosophy of Science* 69(S): S342-S353.
- Glennan, Stuart S. (2005). "Modeling Mechanisms." *Studies in the History and Philosophy of Science Part C* 36(2): 443-464.
- Glennan, Stuart S. (2010a). "Mechanisms, Causes, and the Layered Model of the World." *Philosophy and Phenomenological Research* 81(2) 362-381.
- Glennan, Stuart S. (2010b). "Ephemeral Mechanisms and Historical Explanation." *Erkenntnis* 72(2): 251-266.
- Hempel and Oppenheim (1948) "Studies in the Logic of Explanation" *Philosophy of Science* 15(2), 135-175.
- Hempel, Carl (1963) *Explanation and Prediction by Covering Laws* Philosophy of Science: The Delaware Seminar, New York: Wiley.
- Illari, Phyllis, and Jon Williamson (2010). "Mechanisms are real and local." in *Causality in the Sciences*, eds. Illari, Russo, and Williamson. Oxford University Press.
- Machamer, Peter; Darden, Lindley; and Craver, Carl F. (2000). "Thinking About Mechanisms." *Philosophy of Science* 67(1): 1-25.
- Mitchell, Sandra D. (2008). "Exporting Causal Knowledge in Evolutionary and Developmental Biology." *Philosophy of Science* 75(5): 697-706.

- Mitchell, Sandra D. (1997). "Pragmatic Laws." *Philosophy of Science* (special issue) PSA 1996: Part II, Symposia Papers, ed. Darden: S468-S479.
- Perini, Laura (2005a). "The Truth in Pictures." *Philosophy of Science* 72 (1): 262-285.
- Perini, Laura (2005b). "Visual Representations and Confirmation." *Philosophy of Science* 72 (5): 913-926.
- Salmon, Wesley (1978). "Why ask, 'Why?' An inquiry concerning scientific explanation." Proceedings and Addresses of the American Philosophical Association, 51, 683-705.
- Salmon, Wesley (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton University Press.
- Strevens, Michael (2008) *Depth: An account of scientific explanation*. Harvard University Press.
- Trout, J.D. (2007). "The Psychology of Scientific Explanation." *Philosophy Compass* 2 (3): 564-591.
- Williamson, Jon (2011). "Mechanistic Theories of Causality Part I." *Philosophy Compass* 6 (6): 421-432.
- Wimsatt, William C. (1994). "The Ontology of Complex Systems: Levels of Organization, Perspectives, and Causal Thickets." *Canadian Journal of Philosophy* 20(S): 207-274.
- Wright, Cory (2012). "Mechanistic explanation without the ontic conception." *European Journal for Philosophy of Science* Springer Online First, DOI: 10.1007/s13194-012-0048-8.