Back to the Cradle: Mechanism Schemata from Piaget to DNA

Catherine Stinson

September 15, 2016

Abstract

Mechanism schemata are one of the least understood parts of MDC's account of mechanistic explanation. Relatedly, there is a common misconception that there is no place for abstraction in MDC mechanisms. These two problems can be remedied by looking more carefully at what MDC say both in their 2000 paper and elsewhere about schemata and abstraction, and by following up on a comment of Machamer's indicating that Piaget was the inspiration for schemata. Darden's work on mechanism discovery reveals an important role for abstract mechanism schemata, not only as concise representations of known mechanisms, but also as hypotheses for how unknown mechanisms might work. Piaget's schemata likewise both capture the developing child's current state of knowledge, and serve as templates into which new experiences are fit. If we assume that mechanistic explanations should pick out real difference-makers, and that what makes a difference needn't always be fine-grained details, we open up an explanatory role for abstract schemata to play. Schemata represent the coarser-grained features of mechanisms that are often among the most important difference-makers. In many contexts they might need to be supplemented with finer grained details in order to give an adequate explanation, but independent of those details, schemata remain explanatory in virtue of being part of the explanation. Detailed accounts of real difference-makers need not be seen as in competition with abstract models or generalizations. They can (and should) be combined in mechanistic explanations.

1 Introduction

While writing my dissertation under the joint supervision of Peter Machamer and Ken Schaffner, there were several occasions when Peter and I discovered that we had diametrically opposed views. Once or twice Ken quietly assured me after the end of the meeting that he agreed with me. But more often I would spend the next several months or years trying to figure out how Peter could possibly be right about something as obviously wrong as, for example, top-down causation.

One thing that we largely agreed on was the role of abstraction in mechanisms, although we seemed to disagree with nearly everyone else. We did, however, disagree on how to pluralize "schema," which he told me had also been a bone of contention during the writing of Machamer et al. (2000) [hereafter MDC] too. Peter was for 'schemata.' Lindley Darden was for 'schemas.' Peter prevailed. I suspect the disagreement over spelling belies a deeper ambiguity, which is manifest in MDC's unclarity about what a schema is.[1]

At this point in the conversation Peter switched from his usual booming voice to the softer one he uses when he's sharing deeply held beliefs, to tell me that what he had in mind for schemata was Piaget. I nodded in silent understanding, although I had no idea what he was talking about. Several years later I've finally gotten around to looking it up. Here I will use Piaget's schemata to help clarify some misunderstandings of mechanism schemata, and to show how abstraction might fit into the MDC view of mechanistic explanation.

I begin by looking at how schemata are described in MDC, how they are understood by subsequent commentators, and how they are described in Darden's work. Next I introduce Piaget's schemata and draw out the characteristics of those that Peter had in mind, in order to elaborate how mechanism schemata are meant to capture abstract difference-makers with broad applicability. I contrast that with the popular reception of the MDC view, where Craver in particular is taken to speak against there being an explanatory role for abstract mechanisms. Craver's focus on detail in causal mechanisms seems to preclude the essential role that schemata play not only in mechanism discovery, but also in explanation. I then describe a way of interpreting MDC that makes room for both detail and abstraction, and that is more consistent with the textual evidence than the popular interpretation.

2 Mechanism Schemata

The MDC characterization of mechanism says, "Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions" (Machamer et al., 2000, 3). Part of the paper that has Peter's stamp on it is the focus on types of mechanisms: "Scientists... are typically interested in types of mechanisms, not all the details needed to describe a specific instance of a mechanism. We introduce the term 'mechanism schema' for an abstract description of a type of mechanism" (Machamer et al., 2000, 15). In this passage, schemata sound something like templates or abstract models, and it is noted that the mechanisms scientists are interested in are typically not highly detailed.

The very next sentence is more often quoted though: "A *mechanism schema* is a truncated abstract description of a mechanism that can be filled with descriptions of known component parts and activities" (Machamer et al., 2000, 15). This is often paired with with the definition of a mechanism sketch: "For

[1] This makes one wonder whether there might be deep messages hidden in some of Peter's other idiosyncratic spelling preferences.

epistemic purposes, a *mechanism sketch* may be contrasted with a schema. A sketch is an abstraction for which bottom out entities and activities cannot (yet) be supplied or which contains gaps in its stages" (Machamer et al., 2000, 18). The contrast between mechanism schemata having known components while the entities of a sketch are not yet known suggests that the difference is just a matter of the state of our knowledge about the mechanism. The case study MDC present of discovering the mechanisms of protein synthesis places the schema of the central dogma of molecular biology as a later stage of mechanism discovery after Watson's earlier sketch, which may also lend this impression. Craver later describes schemata in just this way. He says, "Between sketches and complete descriptions lies a continuum of *mechanism schemata* whose working is only partially understood" (Craver, 2007, 114). This is one of the only passages in Craver's book where he mentions schemata at all. Many commentators have understandably followed Craver in interpreting schemata as intermediate between sketches and fully characterized mechanisms.

I think there is a better interpretation of MDC's schemata, which is more consistent with how Darden and Machamer describe schemata in other places. The contrast between sketch and schema in terms of our state of knowledge can be made "for epistemic purposes," but the inclusion of that preamble hints that scientists might also have other purposes. Darden makes the epistemic distinction between models of mechanisms at different stages of development in terms of a box metaphor: "The goal in mechanism discovery is to transform black boxes (components and their functions unknown) to gray boxes (component functions specified) to glass boxes (components supported by good evidence), to use Hanson's (1963) metaphor. A schema consists of glass boxes; one can look inside and see all the relevant parts" (Darden, 2008). But Darden also describes another, arguably more important, job for schemata to do: "Schema instantiation provides an abstract type of mechanism that may be specified to apply to a particular case. Types of mechanisms may be depicted in abstract mechanism schemata; instantiation is the process of making a schema less abstract and applicable to a particular case" (Darden, 2002, S355). This is one of three reasoning strategies she describes for discovering mechanisms.

Darden describes how analogous theories can be grouped together and an abstract schema can be constructed that covers them all. Wave phenomena and natural selection are two examples she mentions. She does not go into much detail here, but presumably the reason for describing types of mechanisms is that they have shared features, which can be explained in the same way. What we know about one kind of wave might turn out to be true of waves in general. This allows us to apply what we learn from experiments in water basins to electromagnetic radiation, for example. Furthermore, each new mechanism we discover provides a new way of approaching phenomena for which we as yet have no good explanation. Darden cites the example of how the discovery of reverse transcriptase "opened up a space of possible mechanisms with feedback into DNA from elsewhere," such as proposals for directed mutation, and feedback from the soma to the immune system (Darden, 2002, S361).

What I take from this is that once we've described a mechanism adequately in one context, we might represent it abstractly as a schema for various purposes. As a representation of the original mechanism, it is a glass box that we can look through to see the details. This might serve many practical and cognitive purposes such as communicative ones. As a representation of a type of mechanism, it can also serve epistemic, explanatory and theoretical purposes; we can use it to transfer knowledge about one instance of the type to others (or the type in general), or to form a hypothesis for how a phenomenon we don't yet understand might work. In these latter cases, where the schema is used as a template that can be applied to many contexts, the schema no longer acts as a glass box. Depending on how much we know about the new context, we might be able to fill in the new details right away and immediately have an adequate explanation. Much more likely is that we will be unsure of some of the relevant details, or unsure of the hypothesis that the borrowed schema applies to this case. Here what was a schema of the original mechanism becomes a sketch of the new one. We need to do some work to fill in the gaps and evaluate whether the schema-cum-sketch is helpful.

3 Piaget's Schemata

Piaget's schemata work in very similar ways. His theory of how children develop intelligence goes roughly as follows: Intelligence is the result of a complex interaction between the child's environment and their reasoning abilities (which are partly biologically determined, partly a result of this ongoing interaction). Through a series of experiments that Piaget performed on his children (with the help of their nurse), he observed how sucking and rooting instincts gradually develop into coordinated nipple-finding skills, how undirected looking develops into visual acuity, and so on.

The sucking schema, in Piaget's simplest example, begins as a reflex to suck whenever something touches the insides of the lips. The infant will try sucking blankets, fingers, or other body parts if it comes into contact with those. The schema gradually becomes more complex with experience. New things the baby encounters, such as its own thumb, might get added to the schema, and its movements become more coordinated such that it can automatically turn its head in the direction of the nipple, and move its thumb to its mouth. This requires a process Piaget calls *assimilation*, which involves cumulative repetition, generalization (incorporation of new objects), and finally motor recognition (Piaget, 1952, 36-37). A second process called *accommodation* can alter the schema to better account for experience. When the schema does not lead to satisfactory results, for instance when the infant sucks something it doesn't like, accommodation will ensure that the infant won't try sucking the offending object again. The coordination of two schemata, such as is required for recognizing the nipple by sight, also involves accommodation. These processes of assimilation and accommodation working together in equilibrium are what Piaget calls *adaptation*.

In Peter's analogy, "understanding the world through Piaget's schemata is like understanding the world through or by mechanisms" (personal communication). We can think of the mind of the child as being like the state of scientific knowledge, and the schemata through which the child makes sense of their world as akin to the mechanism schemata through which scientists make sense of nature. Both start out by fumbling around taking in information indiscriminately, gradually collecting experiences, grouping them together, and drawing associations between them, which coalesce into hypotheses about how the world works, or schemata for understanding it. Both continue gathering information, taking note of exceptions, and making alterations to the schemata when surprises crop up. Paraphrasing Peter, we can think of some schemata (causal schemata in particular) as ideas of mechanisms. He says, "Sometimes the schema you need to understand something is the mechanism by which it works... And since schemata are always modifiable with new input (accommodation) they are always in a sense preliminary and in need of development" (ibid).

This is quite a bit different than the popular interpretation of mechanism schemata as glass box descriptions of well understood mechanisms, or later stages in the development of a sketch. Sketches typically are used to represent the (incomplete) state of our knowledge. They only apply to one phenomenon, and usually only increase in epistemic support. Schemata, on the other hand, are not primarily employed as representations, and part of their purpose is to be widely applicable. An emphasis I would add is that the primary purpose of schemata (both Piagetian and mechanist) is that they can be used to make sense of experience. Piagetian schemata are supposed to help us understand new experiences, so applicability to various cases is their very purpose. The more we can understand with a given schema, the better a schema it is.[2]

Schemata are formed based on what we have already experienced, so in a sense they represent that past experience, but the purpose of them is to help us understand what we have yet to experience (or what we were not yet able to understand when we did experience it). Going back to Darden's point, we develop mechanism schemata (at least in part) in order to be able to instantiate them in novel ways. A mechanism schema might start out as an abstract description of a particular phenomenon, but its main purpose is to help us connect that phenomenon to other things. These connections may be to phenomena we already understand, or to new phenomena that we don't yet know how to understand. They are more like working hypotheses, with the emphasis on *working*. Schemata may fall anywhere on the spectrum of epistemic support from complete guesses to well-supported theories, and the same schema might have a very different epistemic status when applied to different contexts.

Examples of this use of mechanism schemata are as common as analogical reasoning in science, which is to say utterly ubiquitous. An example I like is Hinton's 'wake-sleep' algorithm. The 'wake-sleep' algorithm is a method for unsupervised training in neural networks, i.e., a way for a vaguely brain-like artificial

[2] I do not mean to say that generality is the only or even the main characteristic that we look for in schemata. A schema that can make sense of a previously mysterious experience is also valuable, even if that schema is only narrowly applicable. Bogen (2005) defends a similar point.

intelligence system to teach itself how to improve its performance without the need for a teacher or a set of training data to explicitly tell it what it did wrong or right. One way in which a schema from elsewhere is used in this algorithm is apparent from the name; the training method was inspired by psychological theories about the function of dreaming. The algorithm cycles between 'wake' and 'sleep' phases. In the 'sleep' phase, the network generates stochastic 'fantasy' vectors at the output layer. "These fantasies provide an unbiased sample of the network's generative model of the world" (Hinton et al., 1995). In other words, the fantasies represent the network's encoding of its inputs up to that point in training. Based on these fantasies, the network's parameter settings are adjusted such as to maximize the probability of recreating the input that could have produced that fantasy as output. Metaphorically speaking, the network tries to make sense of its thoughts by changing the very machinery that gives rise to them in such a way as to make those thoughts unsurprising.[3] Furthermore, the method Hinton uses for adjusting the parameter settings is a schema borrowed from physics. The method is "precisely analogous to the way in which the energies of the alternative states of a physical system are combined to yield the Helmholtz free energy of the system" (Hinton et al., 1995).

Hopefully the idea is clear despite this being a rather technical example. A hypothesis about how dreaming works, and a bit of math developed in an unrelated branch of physics were redeployed here to do different jobs than they were originally developed to do. Instead of starting from scratch and building a neural network training method from the ground up, Hinton borrows ideas that he knows already work to solve new problems. To put it bluntly, the kind of science that gets published in *Science* creatively adapts successful ideas from other fields, abstracts away the details, then re-applies them to new cases. Mechanism schemata are the abstract models of mechanisms that can be so deployed.

4 MDC's Anti-abstraction Reputation

Darden and Machamer's view(s) of schemata highlight these ways in which abstraction is useful in mechanistic explanation. Curiously, several recent papers which also argue for the neglected role of abstraction in mechanistic explanation cast MDC as their opponents.

Kuorikoski argues that there are two distinct concepts of mechanism: "First there is the concept of mechanism as a componential causal system, which is accompanied with the heuristics of decomposition and localization. Second, there is the concept of mechanism as an abstract form of interaction, accompanied by the strategy of abstraction and simple models" (Kuorikoski, 2009). In cases from economics, he shows that it is relational properties between entities rather than properties of their parts that do the explaining. He goes on to say that, "Since the form of interaction is not in itself dependent on the way the causally

[3] Incidentally, this is not unlike Piaget's account of the development of intelligence.

relevant properties of the component parts are constituted..., the same sample models and hence 'the same' mechanism schemata can be utilized in many different kinds of contexts or domains" (Kuorikoski, 2009, 152). The problem Kuorikoski raises is not just that we seem to want to use the term 'mechanism' for both token causal chains and for abstract types of interactions, but furthermore, that in many branches of science, the factors that seem to do the work in explanations are not the details about the parts, but rather relational properties between these parts. He complains that recent accounts of mechanism are primarily concerned with the componential causal system concept, and calls out MDC as being "especially" guilty in this regard (Kuorikoski, 2009, 147).

Levy and Bechtel (2013) also argue for the value of abstraction in explanation, focusing on the Hodgkin-Huxley model. They complain that accounts of mechanistic explanation miss cases where "a model is *deliberately* 'sketchy', i.e. where gaps aren't the product of ignorance or theoretical limitations, but of an intentional strategy" (Levy and Bechtel, 2013). This sounds like what Peter had in mind for mechanism schemata, however, MDC are named as a set of authors whose line of thinking requires correction because they tend to "emphasize completeness and specificity" (Levy and Bechtel, 2013). The cases Levy and Bechtel cite as benefiting from more abstract explanations are ones where lower-level entities are treated as collections or aggregates, such that the individual details of each entity are less important to the explanation than the properties of the collective. This is of course just one way in which abstraction from details can be useful. There are many other abstract types that schemata might capture beyond just aggregates, as Darden's examples illustrate.

Since Craver has been the most prolific of the three in expanding on the MDC view of mechanism, and his 2007 comments about schemata seem to take a different direction, it is understandable that the role of schemata as abstract mechanisms in the MDC account has largely escaped notice. In several places (Craver, 2006, 2007; Kaplan and Craver, 2011) Craver even seems to suggest that explanatory adequacy cannot be achieved with abstract models. He claims that "merely subsuming a phenomenon under a set of generalizations or an abstract model" does not explain the phenomenon (Craver, 2006). He warns against models that provide an equation that describes the phenomenon, but without specifying what the underlying mechanisms are that make the equation a good fit. He criticizes how-possibly models, or "loosely constrained conjectures about the mechanism that produces the explanandum phenomenon" (Craver, 2006), and argues that Hodgkin and Huxley's equations describing ion flow during the action potential were not truly explanatory until correspondences between the equation's parameters and entities acting in cell membranes were confirmed (Craver, 2008).

I take it that these claims are concerned not so much with whether abstract descriptions of mechanisms play any explanatory role at all, as with whether abstract descriptions are sufficient to explain. Craver's point is perhaps not that abstract models have no explanatory value, but rather that for descriptions of any kind to be explanatory they must refer to what's really making the explanandum happen. Craver is often interpreted as arguing for a 'more details better' account of mechanistic explanation, according to which mechanisms must include all the gory details that could under any circumstances be relevant to their operation in order to be explanatory.[4] I take Craver's point about detail to be a (not entirely convincing) account of how to draw boundaries around mechanisms, not about how much detail is required in any particular mechanistic explanation. In epistemic accounts of explanation it is all too easy to find a role for abstraction: abstract descriptions, metaphors, and analogies are all excellent ways of communicating information about mechanisms, regardless of whether anything in the world literally corresponds to them. MDC's explanatory concerns are focused on the mechanisms themselves, however, and their productive continuity. Craver later recasts productive continuity in terms of Woodward's (2003) interventionist account of causation. The MDC account is not anti-abstraction; it just requires that abstract models refer to the actual difference-makers in order to be considered explanatory.

5 Abstract Schemata and Explanation

What remains is to show how abstract mechanism schemata can act as explanations, which for MDC implies that they refer to the actual difference-makers. It should be uncontroversial that at least sometimes we need the gory details to explain how a mechanism works. To see why we needn't always go into all the details, and how details and abstraction might be combined, it is helpful to compare the 'more details better' argument attributed to Craver to Kim's argument against non-reductive physicalism (Kim, 2005).

In Kim's argument, the possibility that there might be autonomous mental causation acting at a higher level is rejected, because the lower-level physical events they supervene on also seem to be causes of subsequent events, and causes shouldn't be overdetermined (by assumption). Just as Kim chooses to have mental causes be reducible to physical ones in order to avoid his dilemma, the 'more details better' interpretation of Craver locates the real difference-makers in the gory details. Where the comparison breaks down is that an essential assumption in Kim's argument is causal closure of the physical domain. In the MDC account, both higher and lower-level mechanisms are physical, so causal closure does not block the possibility that a lower-level event could be caused by a higher-level one. The physical cause of the lower-level event has to also be physical, according to causal closure, but needn't be at the same mechanistic level. Causal relationships can cross-cut mechanistic levels.

Machamer and Sullivan (2001) question the possibility of unabmiguously identifying mechanistic levels because of the ubiquity of cross-level relationships. There are cases where causal relationships cross levels, and cases where particular entities can appear in multiple levels alongside other entities of different scales. As I argued in Stinson (2016), there is little reason to believe that mechanisms must be organized into neat hierarchies following mereological

[4]See Craver's unpublished manuscript "Are More Details Better?" for his most recent take on this.

relationships. Instead of talking about levels of mechanisms, I'll talk in terms of a continuum from fine-grained detail to abstraction.[5]

This rejection of an analogue to Kim's causal closure argument, and the recognition that mechanistic levels may cross-cut each other both in terms of composition and causal relationships suggests another way to solve the puzzle. Certainly some of the causes at play in a mechanism might be fine-grained ones. These would typically be responsible for small effects that would only be relevant in some explanatory contexts. But other causes at play in a mechanism might be coarser-grained ones. These would typically be responsible for the major effects and would be relevant in most explanatory contexts. (Of course small details sometimes have large effects, and vice versa.) One of the reasons why mechanism schemata are so useful is that they pick out the major difference-makers, so focus attention on the features of the mechanism that are most important.

As described above, schemata start out as glass box descriptions. These are abstract representations of mechanisms that we already understand. Some details are removed in order to show more clearly how the mechanism works. Schemata are often assumed to have the potential to be applied more widely. They can act as hypotheses for how unknown mechanisms may work in what Darden calls schema instantiation. Here they are turned into sketches. These sketches need to be filled in and empirically justified in order to gain the status of explanations. If in the new context the schema-cum-sketch turns out not to correspond to the underlying mechanisms, then it is not an explanation. This is the sort of case Craver warns against. However, when the sketch does prove to have been successful, the original schema explains the new phenomenon too. (It may also happen that the original schema gets altered to fit the new case, which may result in the old explanation being adjusted too. This can be understood in terms of Piaget's assimilation and accommodation.) When a more abstract model can serve to unify the explanations of several particulars, this is usually thought of as an explanatory virtue. Although I wouldn't endorse the view that unification is *the* goal of explanation, it is one among many explanatory virtues. Of course an abstract schema will rarely be a complete explanation of either the original or the new phenomenon. Some of the details of each case might also be required in a complete or even adequate explanation. Most adequate explanations are probably combinations of multiple models at various levels of abstraction. Different combinations of these multiple models would be called for in different explanatory contexts. Nevertheless, the abstract schema is an explanatory model in virtue of being part of the explanation.

To give a very simple example, a bag filled with sand used as a draught stopper at the base of a door does its job of stopping draughts in virtue of it being heavy, dense, and heat resistant enough not to let cold winds blow in. Each grain of sand in it does part of the job, but it's the fact that it's a lot

[5]The continuum of granularity that Weiskopf (2011) describes refers to "size of the chunks into which one decomposes a mechanism." My continuum of abstraction is closely related, but refers to how much detail we include in our descriptions, regardless of the size of the components. Both small and large things can be described in more or less detail.

of them bound together in a bag that matters most to it serving its purpose. This is a case of an aggregate, like the ones Levy and Bechtel discuss. In some other contexts, the finer-grained properties of the sand might matter. If we wanted to know how the sandbag's draught stopping power changes after being hit by lightning (an exceptional case of the type Craver might worry about), it would be relevant to know what proportion of silicon dioxide and calcium carbonate the sandbag contains. If the particular shape of the glass created by the lightning strike also affected the sandbag's draught stopping properties, it might be necessary to know the distribution of grains of particular chemical compositions in different parts of the sandbag. In each case we have the same bag of sand, but what matters about that bag of sand shifts depending on the explanatory context.

Regardless of how much detail we might need, a complete description that includes all the causally-relevant aspects of the mechanism needn't be a collection of just fine-grained entities and activities like an exhaustive description of every grain of sand. It would instead be a collection of models at various levels of grain and abstraction. It might include things like the mass, shape, and wind permeability of the whole bag, and the distributions of the different particle types in the bag, in addition to the locations of individual grains. We can leave the finer-grained details out of the explanation in cases where the locations of individual grains don't make a difference. But we also want to maintain an explanation that corresponds to the actual difference-makers. In these cases, the explanation might just include the mass and shape, or the distributions of grain types. Abstract explanations like these still refer to the actual difference-makers.

One minor quibble is that it is not clear here whether what we call 'the mechanism' is what is picked out by the complete description regardless of the explanatory context, or whether what constitutes the mechanism shifts as the context changes to only include the components that are causally relevant in that context. In one sense these pick out the same thing in the world. The sandbag as a massive object of a particular shape doesn't leave out any of the grains. We can use schemata to pick out the relevant aspects of the mechanism and abstract away the details. This abstracting away doesn't lose details the way one might lose grains of sand if they were to spill out a hole in the sandbag. Schemata still capture the entire mechanism, just under a coarser-grained description.

Schemata aren't the loosely constrained conjectures about mechanisms that Craver warns against. They are not just less sketchy sketches. They have a different epistemic, explanatory, and theoretical status. One might know in great detail about a mechanism and the related components and sub-mechanisms that surround it, but nevertheless have good reason to think that in some contexts what makes a difference is just a coarse-grained aspect of that cluster of components. It might only be a mechanism's information-processing properties that make a difference to what it does, and not how those are implemented. It might only be the fact that it obeys the Hodgkin-Huxley equations that makes a difference, and not the finer details of exactly what makes it obey those equations. The most important difference-makers are often abstract features of the system, as Kuorikoski and others have pointed out. These abstract features of mechanisms are what schemata describe, and as MDC say, this is what scientists are typically interested in.

Take the somewhat more complex example of a unicycle. Insofar as the unicycle is an instance of a general gyroscope schema, its angular momentum is conserved and it stays upright. It works much the same way as the gyroscopes in helicopter navigation systems and iphones. Insofar as its wheel is out of true, it wobbles. Although there might be some explanatory contexts where finer details about the wheel are needed, the more general facts are not jettisoned; they are added to. If the ground is smooth, angular momentum might be enough to explain why the unicyclist stays upright. If the ground is bumpy, the fact that the wheel is out of true might also be required for an adequate explanation.

The role that abstract mechanism schemata play in discovery carries over to explanation. What mechanism schemata do is highlight the features of a mechanism that are the most important difference-makers, and show how they work in general. When we abstract the details away from a mechanism to create a schema, it should pick out these important features. If we then instantiate the schema in a new context, the hope is that the features highlighted in the schema work in roughly the same way as the features playing the same role in the original mechanism. If the instantiation is successful, we assimilate the new case into our schema. We might also find out something new about the schema in the process of trying to instantiate it, and alter the schema to accommodate that new knowledge. In this sense schemata are always open to change, as is all of our scientific knowledge.

6 Conclusion

I began with a disagreement over how to pluralize 'schema' and a mysterious remark about Piaget. I drew together M, D, and C's various comments on schemata from other texts to get a more accurate picture of their view(s) than the popular reading. By combining Darden's work on schema instantiation in mechanism discovery with Peter's Piagetian inspiration, I constructed a picture of mechanism schemata as abstract templates that are formed by abstracting away details of experience, and that are used to help understand new experiences. Craver's focus on details and insistence that explanations reflect the actual difference-makers are often taken as precluding any explanatory role for abstract models or generalizations. I showed how abstract models like mechanisms schemata can reflect the actual difference-makers, and be combined harmoniously with more detailed models. The role of schemata in mechanism discovery carries over to explanation. Because of how they are constructed through a process of ongoing assimilation and accommodation, schemata are uniquely placed to serve as hypothesized explanations. Once they receive empirical support in new contexts, mechanism schemata become genuine, albeit partial, explanations.

References

Bogen, J. (2005). Regularities and Causality: Generalizations and Causal Explanations. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36(2):397–420.

Craver, C. F. (2006). When Mechanistic Models Explain. *Synthese*, 153(3):355–376.

Craver, C. F. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford University Press.

Craver, C. F. (2008). Physical Law and Mechanistic Explanation in the Hodgkin and Huxley Model of the Action Potential. *Philosophy of Science*, 75(5):1022–1033.

Darden, L. (2002). Strategies for Discovering Mechanisms: Schema Instantiation, Modular Subassembly, Forward/Backward Chaining. *Philosophy of Science*, 69(3):S354–S365.

Darden, L. (2008). Thinking Again about Biological Mechanisms. *Philosophy of Science*, 75(5):958–969.

Hinton, G. E., Dayan, P., Frey, B. J., and Neal, R. M. (1995). The "Wake-Sleep" Algorithm for Unsupervised Neural Networks. *Science*, 268(5214):1158–1161.

Kaplan, D. M. and Craver, C. F. (2011). The Explanatory Force of Dynamical and Mathematical Models in Neuroscience: A Mechanistic Perspective. *Philosophy of Science*, 78(4):601–627.

Kim, J. (2005). *Physicalism, Or Something Near Enough*. Princeton University Press.

Kuorikoski, J. (2009). Two Concepts of Mechanism: Componential Causal System and Abstract Form of Interaction. *International Studies in the Philosophy of Science*, 23(2):143–160.

Levy, A. and Bechtel, W. (2013). Abstraction and the organization of mechanisms. *Philosophy of Science*, 80(2):pp. 241–261.

Machamer, P., Darden, L., and Craver, C. F. (2000). Thinking about Mechanisms. *Philosophy of Science*, 67(1):1–25.

Machamer, P. K. and Sullivan, J. A. (2001). Leveling Reductionism. http://philsci-archive.pitt.edu/id/eprint/386.

Piaget, J. (1952). *The Origins of Intelligence in Children*. International Universities Press, New York.

Stinson, C. (2016). Mechanisms in Psychology: Ripping Nature at its Seams. *Synthese*, 193(5):1585–1614.

Weiskopf, D. A. (2011). Models and Mechanisms in Psychological Explanation. *Synthese*, 183(3):313–338.

Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation.* Oxford University Press, USA.