Daniel Sirtes,
Science Studies Program/Philosophy Department, University of Basel, Switzerland
For comments, questions, and citation requests please contact me:
daniel.sirtes@unibas.ch

A Pragmatic-Ontic Account of Mechanistic Explanation

Craver's (2007) account of explanation in neurobiology offers one of the most sophisticated explications of the mechanism concept. This paper argues that despite groundbreaking advances in understanding mechanistic explanation, serious challenges remain. The first goal of this paper is to address the notorious problem of explanatory relevance concerning mechanistic explanation. I argue that Craver underestimates the importance of pragmatic constraints on the individuation of mechanisms, and that his suggestion for a solution of the explanatory relevance problem is therefore insufficient on several counts. My second goal is to develop an alternative that explicitly incorporates both pragmatic and ontic aspects of mechanism individuation.

1. Introduction

Although (Bechtel and Richardson 1993) and (Glennan 1996) re-awakened the concept of a mechanism in contemporary discussions of explanation from a slumber of about ten years,[1] the new mechanism debates were largely initiated by (Machamer, Darden, and Craver 2000) (hereafter, MDC).[2] MDC claim that mechanisms are sought to explain how a phenomenon comes about or how some significant process works. They define mechanisms as "entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions" (MDC, 3).[3] They believe that thinking in terms of mechanisms provides a new framework for addressing many traditional philosophical issues such as "causality, laws, explanation, reduction, and scientific change" (2000, 1). This is the 'Mechanista Manifesto' that has motivated a remarkable surge in publications on this topic.[4] The vast majority of the articles discussing MDC are either attempts to develop their account further, or attempts to apply the very same concept of mechanism to different disciplines and to different philosophical issues. The mechanistic view has been extended to disciplines ranging from computer science through chemistry to sociology (see e.g. (Piccinini 2007), (Thagard 2003), (Mayntz 2004), respectively); and to philosophical issues including reductionism, feminism

---

[1] The concept had its last prominent appearances in Wimsatt (1976), Railton (1980), and Salmon (1984).

[2] MDC is now the most cited paper in *Philosophy of Science* according to its own website.

[3] Machamer claims in a later paper that 'regular' should be dropped from the definition (Machamer 2004, 37) agreeing with Jim Bogen's arguments that there might be mechanisms that operate only sporadically, or even just once (published later in (Bogen 2005).

[4] I thank Ken Waters for introducing me to this slightly funny, but spot on, name for the fervent supporters of this position.

and objectivity (e.g. (Delehanty 2005), (Fehr 2004), (Thagard 2008), respectively). By contrast, there have been relatively few attempts to examine critically the account itself, or its relation to more traditional accounts of explanation and causation.[5] The aim of this paper is to do exactly that. I begin by examining a precursor to the recent mechanism movement, the causal-mechanical account of explanation, in order to introduce a critical distinction that accounts of explanation should address (especially if they are ontic). The distinction is between *ideal, infinite* vs. *pragmatic, partial* descriptions of the explanans. This points to a traditional challenge that any account of explanation must tackle: the problem of explanatory relevance (Section Two). The traditional problem of explanatory relevance for the deductive-nomological account concerns the inclusion criteria for the premises of an explanation in order for them to be necessary and sufficient for the derivation of the explanandum. In Section Three, I illustrate how this problem finds its counterpart in mechanistic explanation. The description of the mechanism, whether it is ideal or partial, should include those and only those entities and causal interactions that are explanatory of the explanandum phenomenon. I argue that only a pragmatic account of mechanistic explanations can deliver a satisfactory solution for this problem. An ideal description would be extremely redundant and so complex that it cannot be a viable characterization of explanation in science. Section Four takes a closer look at Craver's account of explanandum phenomena. I show that he vacillates between ideal and pragmatic notions of explanation resulting in an inconsistent account. In Section Five, I challenge Craver's own take on the explanatory relevance criteria (for constitutive explanation), called 'mutual manipulability', which will turn out to be insufficient on several counts. In Section Six, drawing on van Fraassen's pragmatic theory of

[5] Notable exceptions are (Weber 2008) and (Psillos 2004).

explanation and the Woodward/Craver criteria of manipulability, I suggest an alternative account of mechanistic explanation that, together with a causal-mechanical determinant of the relevance relation, explicitly incorporates the unavoidable pragmatic constraints on the individuation of mechanisms. I illustrate this in Section Seven with an example from molecular biology, the mechanism of transcription. I conclude by arguing that mechanisms are not “acting entities” or even kinds of things; as every appropriately posed question seeking a mechanistic explanation will lead to an answer that includes a description of a different mechanism. Thus, Mechanisms are best thought of as any set of entities and causal interactions that are referred to in such an explanation.[6]

2. Back to the roots: The ideal-pragmatic distinction in ontic explanations

Both Peter Railton and (*post*-1984) Wesley Salmon promoted a causal-mechanical view of explanation. They both suggest that to explain something is to show how the explanandum fits into the causal nexus of the world. This causal nexus is the ontic basis for explanation, as it constitutes what there is in the world. It is composed of myriads of causal processes and causal interactions. To explain is to show how a phenomenon fits into this nexus by describing the processes and interactions that lead to that phenomenon.[7] The concept of ‘mechanism’ plays a relatively minor role in both Salmon and Railton’s accounts. Railton claims that he does

[6] The question how mechanistic questions are posed *appropriately* is of course part and parcel of a pragmatic-ontic account of mechanistic explanation.

[7] Railton’s account differs from Salmon's in a number of significant respects, and it can be argued that Salmon distorted Railton’s *opus magnum* ‘Explaining Explanation’

> not have anything very definite to say about what would count as ‘elucidating the mechanisms at work’ - probabilistic or otherwise - but it seems clear enough that an account of scientific explanation seeking fidelity to scientific explanatory practice should recognize that part of scientific ideals of explanation and understanding is a description of the mechanisms at work, where this includes, but is not merely, an invocation of the relevant laws (Railton 1978, 242).

Salmon only states that “Causal processes, causal interactions, and causal laws provide the mechanisms by which the world works; to understand *why* certain things happen, we need to see *how* they are produced by these mechanisms” (Salmon 1984, 132). Neither are interested in an elaborate explication of the concept of *a mechanism* but rather of mechanistic explanation, and I suspect that the problems I will address may have been a (tacit) reason to steer away from this quagmire.

Nevertheless, a distinction Railton introduced will be helpful in pinpointing the problems with Craver’s (2007) account and its difference to the pragmatic-ontic alternative that I develop. Railton distinguished two basically different conceptions of explanation: *ideal* and *salient* explanations. An *ideal explanatory text* is a potentially infinite text that includes all possible explanations of a phenomenon, including causal, reductive, etiological (i.e. the entire causal history of an event) and structural explanations, on all possible descriptive levels (microphysical, phenomenological, etc.). By contrast, *partial explanatory* information includes only those *salient* parts of the ideal explanatory text that are necessary for answering a specific why-question.

---

by suggesting that there are more similarities with his own account than there really are.

Whereas an ideal explanatory text should include all *objective,* causally relevant information, partial explanatory information is a *pragmatic* notion that includes all the relevant information in a given explanatory context. Here is how Railton characterizes ideal explanatory texts:

> An ideal text for the explanation of the outcome of a causal process would look something like this: an inter-connected series of law-based accounts of all the nodes and links in the causal network culminating in the explanandum, complete with a fully detailed description of the causal mechanisms involved and theoretical derivations of all the covering laws involved. This full-blown causal account would extend, via various relations of reduction and supervenience, to all levels of analysis, i.e., the ideal text would be closed under relations of causal dependence, reduction, and supervenience. It would be the whole story concerning why the explanandum occurred, relative to a correct theory of the lawful dependencies of the world. Such an ideal causal D-N text would be infinite if time were without beginning or infinitely divisible, and plainly there is no question of ever setting such an ideal text down on paper. (Indeed, if time is continuous, an ideal causal text might have to be non-denumerably infinite — and thus “ideal” in a very strong sense.) But it is clear that a whole range of less-than-ideal proffered explanations could more or less successfully convey information about such an ideal text and so be more or less successful explanations, even if not in D-N form (Railton 1978, 247).

This kind of “fully detailed description of the causal mechanisms”, as a part of an *ideal explanatory text*, would be an infinite and ideal account. It would include all the causal events in the past light cone of the explanandum with an ideally (infinitesimally) fine-grained description of the properties, entities and causal interactions. Moreover, it would include a multitude of descriptions on different levels "relative to a correct theory of the lawful dependencies of the world". It seems obvious that such a concept of a mechanism or of the conditions of inclusion in such a mechanism description is not what scientists mean when they talk about mechanisms. Moreover, such a strongly idealized concept of a mechanism is helpful neither for understanding the roles mechanisms play in the heuristics of discovery (see e.g. (Darden 2006)), nor in the integration of results from different fields of research. Scientists are interested in only a part of this ideal text.[8] They are interested in what Railton calls *partial explanatory information*. However, selecting just some part of text as explanatorily relevant from the ideal explanatory text is intrinsically a pragmatic choice as it is based on *salience conditions*. Thus, if the description of a mechanism is the product of pragmatic considerations, the referent of this description (i.e. the mechanism) does not exist *as a mechanism* independently of scientists querying for it. It only exists as a mechanism in relation to some *particular* why-question that makes it explanatorily relevant for us.

---

[8] As both (Darden 2006) and (Craver 2007) show, the heuristic and integrative role of mechanisms is very focused. A missing entity or activity (a black box) in a mechanistic sketch is a starting point for research, and the links between different areas of research are integrated in a mosaic by fitting single entities in different mechanisms described by different means. Mechanism descriptions, as published in scientific articles, are always addressing a particular aspect of the explanandum phenomenon.

This is a conclusion that Craver would like to avoid. For Craver, mechanisms are "objective explanations", which are "full-bodied things" and a good and complete description of such a mechanism represents all of it (Craver 2007, 27). Moreover, "the boundaries of mechanisms — what is in the mechanism and what is not—are *fixed* by reference to the phenomenon that the mechanism explains" (2007, 123). I will argue that this concept of mechanism is untenable not least due to its undesirable ontological consequences. It leads to an over-population of the ontological zoo. There would be as many different mechanisms as there are interesting ways to answer why-questions with mechanistic explanations. Moreover, our choices about which why-questions to answer with which mechanistic explanations would determine what really exists. At times, Craver appears to be aware of these unwelcome ontological consequences, but he tries to have it both ways: mechanism descriptions as ideal explanatory texts *and* as partial explanatory information. As we shall see, by equivocating on the issue, Craver falls prey to one of the most notorious problems in the philosophy of explanation: the problem of explanatory relevance.

3. The newest incarnations of the explanatory relevance problem

The challenge is to answer the question: what conditions must the premises or parts of an explanation satisfy so that *only* those relevant to the explanation are included in the explanans. Every account of explanation must provide a solution to this problem.[9] I would even maintain that it is the touchstone of any account, because

[9] Craver argues that "complete explanatory texts are complete because they represent all and only the relevant portions of the causal structure of the world" and "Explanatory texts can be accurate enough and complete enough, depending on the pragmatic context in which the explanation is requested and given. Objective explanations [i.e. mechanisms] are not variable in this way" (2007, 27). However, as I

if these conditions cannot be satisfactorily stated, the account is incomplete and cannot be fruitfully applied.

Arguably, the problem of explanatory relevance brought most traditional accounts to their knees. Many counterexamples to the deductive-nomological and inductive-statistical models of explanation (e.g. the flagpole and the shadow (Bromberger 1966, 92ff.), the hexed salt dissolution (Kyburg 1965, 147), and the men and oral contraceptive example (Salmon 1971, 34)) aim to show that the conditions for the premises for Hempel and Oppenheim's deductive-nomological model of explanation are either too strong or too weak. In an ironic turn of events, Salmon's (1994) conserved quantity account has been challenged on the grounds that it cannot distinguish between processes that are explanatorily relevant and processes that take place at the same time but are not explanatory.[10] The same problem arises for mechanistic accounts of explanation: How should we demarcate the boundaries of a mechanism? How does one determine what to include as part of a mechanism and what to exclude as irrelevant or unexplanatory?

As Railton already argued, for complete etiological mechanisms, there is no end to the details that could be included in the past light cone of the explanandum event. A similar point applies to the mechanism descriptions that have a temporal boundary set by the explanandum phenomenon. These kinds of mechanisms, such as input-output mechanisms or constitutive mechanisms, include the starting point

---

argue, complete (i.e. ideal) explanatory texts are not descriptions of mechanisms as scientists understand them. Mechanisms in science are to what partial texts refer. I will emphasize this point in the following.

[10] See (Hitchcock 1995). Salmon has acknowledged that this is indeed a problem and re-introduced statistical relevance to his account (Salmon 1997).

and end point of the process to be described, so they are not potentially infinite in the same way that etiological mechanisms are (an ideal mechanism that tracks the history of an explanandum event has no natural starting point in time other than perhaps the big bang). However, an implication of Railton's argument is that even an ideal complete description of a time-limited mechanism is infinite as its time is "infinitely divisible". Moreover, the multitude of levels and different degrees of precision of possible descriptions should add to the plausibility of the claim that complete mechanism descriptions are unattainable in practice, even if one is unconvinced by the argument to that conclusion based on the infinite divisibility of time.

As Beatty (1997) has so poignantly argued, biologists usually do not seek to describe events that occur only once in nature.[11] They aspire to discover mechanisms that are abundant. If we were aspiring at complete descriptions of the processes that bring about a certain phenomenon several times, either at different times or in different individuals, or even in different species, then the only way to attain such a completeness would be a *disjunctive* description, as for example, the paths of the ions in the neuron in the case of an action potential are different in any individual event. Consequently, each occurrence of the explanandum would add another infinite description. This cannot be the answer to anything interesting. So even an *ideal explanatory text* of a type of a phenomenon that has several tokens must somehow abstract from the details that are irrelevant to the explanation of the

[11] There are of course some important exceptions to this, especially in evolutionary biology; most prominently, the description of the beginning of self-reproducing life forms is itself a major goal in pre-biotic evolutionary biology, where a satisfactory explanation of this singular event would be groundbreaking.

explanandum, and therefore may vary between different tokens of the explanandum without changing its identity as the same type of phenomenon. This kind of abstraction from variation that leaves the identity of the explanans processes and the explanandum phenomenon untouched (however this is determined) is a pre-condition for generalization, i.e., for explanations of phenomena types that have more than one instance. It may be that an *ideal explanatory text* includes all the ‘rest’, but it remains unclear how to determine what kind of details, properties, or entities will not end up in the infinite text that includes all true theories, laws, and concepts that might be used to explain the phenomenon. If we restrict the ideal text to mechanistic explanation, i.e. to an *ideal mechanistic explanatory text*, we could of course exclude non-mechanistic explanations, like structural or unificatory explanations. As argued above, the resulting text would still be infinite. However, it might still be sensible to ask whether it is possible to account for the conditions of inclusion in such a text. Thus, the *ideal mechanistic explanatory text* should refer only to those entities and causal interactions that are *relevant* for the explanandum phenomenon. I will call this the *ideal mechanistic explanatory relevance problem* (IMERP).

Railton’s partial explanatory information is an answer to a *specific* question about the explanandum phenomenon. The selection of this text from the ideal text is determined by the kind of question asked. Therefore, the relevance criteria for inclusion in the partial text are (implicitly) *included* in the question itself. Thus, *partial mechanistic explanatory information* is a text that refers to entities and causal interactions that should include only the text that is *relevant* for answering that specific question. This would be the *partial mechanistic explanatory relevance problem* (PMERP). Whereas IMERP asks for the proper relationship between the explanans and the explanandum phenomenon, PMERP asks for the proper relationship between the explanans and the question asked. In order to solve IMERP, we need to explicate the notion of an explanandum phenomenon in a satisfactory manner, including e.g. the numerical identity criteria of a phenomenon (what constitutes just *one* phenomenon in contrast to *two* disparate phenomena), which as I will discuss below, is notoriously difficult. The solution to the PMERP does not need to delve into this quagmire. A proper characterization of questions (using basic erotetic logic and adding further determinants that I will explicate below) is a far easier undertaking that is free of the metaphysical burden tied to the concept of explanandum phenomena. Additionally, a partial text is finite and may be written down in publications and textbooks, something scientists really do, whereas ideal texts are only of concern in philosophical discussions. We will return to my proposed solution to PMERP in Section Six, first however, I would like to show why failing to distinguish between IMERP and PMERP leads to an incoherent account of mechanisms. In order to do so, I will further divide the relevance problems into two sub-problems. First, what are explanandum phenomena of mechanisms (IMERP) or mechanistic questions (PMERP), and second, according to what criteria are the entities and causal interactions that appear in the explanans to be selected (in both relevance problems)? I begin with Craver's account(s) of explanandum phenomena.

4. Craver's two concepts of a phenomenon

The new mechanistic philosophers agree that there are no mechanisms *simpliciter.* Mechanisms are always mechanisms *for* a behavior or a phenomenon, and the description of a mechanism is an explanation of an explanandum phenomenon. Even so, Craver offers two incompatible explications of an explanandum phenomenon. According to one, "It is insufficient to characterize the [explanandum] phenomenon only under standard precipitating conditions. A *complete* characterization of the phenomenon requires one to know…" [all the following:] "… its precipitating conditions, manifestations, inhibitory conditions, modulating conditions, nonstandard conditions, and byproducts" (Craver 2007, 126ff, my italics). Thus, according to Craver, a mechanism should account for all the aspects of the explanandum phenomenon (2007, 128). This characterization has two very distinctive features. First, it claims that there is such a thing as a *complete* characterization of a phenomenon (see also fn. 9). I will not belabor the point that there are infinitely many ways to describe a phenomenon. More importantly, Craver claims that a phenomenon is not just one state or one event, but it is a multitude of states or events. All kinds of events, even the absence of an event, are parts of the phenomenon. For example, inhibitions might result in no event at all. If the explanandum phenomenon is, e.g., the delivery of oxygen to the periphery of the organism, then no delivery of oxygen is also part of the phenomenon (given that the inhibition was causally connected to the entities that are causing delivery under other circumstances). Thus, Craver's first characterization is not of *one process in time.*

Later in the book (Craver 2007, 202-211), Craver gives a very different characterization of the explanandum phenomenon; one that accords much better with other characterizations of a mechanism made by MDC (and Glennan). Craver claims that the explanandum phenomenon is to be identified *contrastively*: "The contrastive description of the explanandum effect is helpful … to identify precisely that feature of the world for which a cause is sought" (2007, 204). Craver alludes to Hitchcock's point that answering the question about whether smoking a pack a cigarettes per day causes lung cancer depends on the contrast to what happened before (was the person a non-smoker or even a heavier smoker before smoking a pack a day) (Hitchcock 1996). This, however, means that an explanandum phenomenon with a different contrast class is to be understood as a different phenomenon. Thus, using Craver's notation, S for a system and Ψ_i for S's activity, S's Ψ_1-ing in contrast to Ψ_2-ing is a different phenomenon than S's Ψ_1-ing in contrast to Ψ_3-ing. Craver illustrates this point with the example of freezing water (2007, 205ff.). If the temperature of a bucket of water at sea level drops to -18.6°C, one could ask both "what caused the water to freeze?" or "what caused the water to freeze in exactly t minutes (rather than faster or slower)?" For the first question, the relevant feature of -18.6°C is only that it is below 0°C. For the second question, it is relevant that it is -18.6°C and not, e.g., -15°C. So while the answer to the first question will only cite the property that the temperature fell under 0°C as causally relevant, the second answer will cite the exact temperature, as "[b]eing below 0°C, however, is relevant only to the fact that the water freezes, not to its rate of freezing. The values of the variable are too coarse to identify the relevant differences" (2007, 206).

This is where the clash between these two characterizations occurs. According to the first characterization, to describe a single phenomenon means to describe *all* of its modulating and inhibiting conditions, that is, to describe S's Ψ_1-ing AND Ψ_2-ing AND Ψ_3-ing and to find a mechanism that explains all these behaviors, whereas according to the second explication, we would have two different explanandum phenomena (S's Ψ_1-ing RATHER THAN Ψ_2-ing AND S's Ψ_1-ing RATHER THAN Ψ_3) and two different mechanisms, accordingly. In the second case, a mechanism is used to answer a question of the sort 'why X rather than Y?', which is a description of a mechanism as *partial mechanistic explanatory information*, as the context of the question is crucial for the selection of the causal processes selected to be included in the mechanism. The first case, by contrast, characterizes the mechanistic explanation as an *ideal mechanistic explanatory text*. To describe all the possible modulating conditions of a phenomenon is an infinite task. A single continuous variable yields an infinite amount of values that should be described. As illustrated in the example of the bucket of freezing water, all temperatures and freezing times are part of this description. The mechanism that accounts for a phenomenon in this wide sense is extremely complex and abstract. We would have to include all kinds of conditionals into its description, as Craver wants us to include laboratory conditions. Experimental set-ups must also be included, and then these must be integrated into different kinds of in vivo systems, etc. It seems, however, that there is something deeply wrong about characterizing mechanisms in this wide sense. Although it is true that one discovers a mechanism by tinkering with it, i.e., by looking at its modulating conditions, inhibiting conditions, etc., identifying the experiments themselves as part of the mechanism or as part of the explanandum phenomenon seems awkward at best. For example, the inhibition of sodium channels by tetrodotoxin (TTX) is one way to understand the role of sodium channels in action potentials. However, the TTX/sodium channel causal interaction is not part of the generic action potential mechanism. The second characterization is much closer to the idea of mechanism as used by biologists, and to what Railton calls partial explanatory information. Such a characterization opens the door to a pragmatic solution to PMERP, which is the basis of the alternative account of mechanistic explanation that I offer.

Craver stresses that his “account is thus no more pragmatic than any view of causation or explanation would have to be” (2007, 204). However, Craver overlooks a decisive difference between causal explanations (as construed by e.g. Woodward) and his account of mechanisms. The former does not imply that there is some kind of ‘fence’ that demarcates the objects and causal interactions cited in the explanation. By contrast, the mechanism itself is supposed to have a distinct boundary, as Craver claims that mechanisms are “full-bodied things” and “acting entities”. If the boundaries that demarcate such ‘things’, and the objects it contains, are different for different questions and depend on how a particular why-question is asked, then we would have to accept a very messy ontology. In ordinary language, we often refer to arbitrary or even gerrymandered ‘things’, like the tip of my finger or a spot on my screen. However, science seeks to avoid such a volatile clutter. It aspires at a lean, stable ontology, with as few entities as possible to explain it all. Accepting the abundance of mechanisms to the basic furniture of the world would run counter its purpose.

However even if we settle for a concept of an explanandum phenomenon that is close to MDC, Glennan and Craver’s contrastive explication (and abandon attempting to solve IMERP), we would still have to confront the second part of the explanatory relevance problem: the criteria for inclusion in an explanation.

5. Mechanistic Explanatory Relevance Criteria

Craver proposes a solution to the problem of explanatory relevance (for constitutive mechanisms). He proposes that the following *mutual manipulability criteria* (MM) are sufficient for constitutive relevance:

> (CR1) When [the entity’s activity] Φ is set to the value φ_1 in an ideal intervention, then [the system’s activity] Ψ takes on the value of f(φ_1).[12]

[12] Craver adds two further constraints to CR1 (CR1a & CR1b) in order to exclude compensatory responses (such as recovery, redundancy, and reorganization) (Craver 2007, 156ff.). I will assume these additional conditions as well, without further elaboration.

and

(CR2) if [the system's activity] Ψ is set to the value ψ_1 in an ideal intervention, then [the entity's activity] Φ takes on the value f(ψ_1).

Thus, according to Craver, in order for an entity and its activity φ_i to be included in the mechanism, it must change the system's activity ψ_i if manipulated, *and* it itself must change if the system's activity is manipulated. Craver contends that these conditions are *sufficient* for inclusion in a mechanism. I will argue that CR1 is necessary, but insufficient for inclusion, and that even adding CR2 does not make it sufficient.

First, Craver's mutual manipulability criterion that determines whether a particular entity is to be included in the mechanism is incomplete because it does not specify the *kinds* of entities that should be included in the mechanism. It does not determine on what *level* of description the mechanism should take place. In (constitutive) mechanisms, we run into the problem that some entities, and parts of these entities, could be manipulated with the same result in the change of the system's activity. For example, manipulating a person's *heart* in a certain way will affect the person's behavior. However, manipulating the person's *pacemaker cells* (which are part of the heart) will also yield the same system activity, e.g. dying or anxiety (as an effect of tachycardia). Thus with Craver's manipulability criterion, we would be *obliged* to include the heart *and* the pacemaker cells as entities in the mechanism, and therefore include arguably redundant descriptions. In order to avoid this, a decision about the *level* of the mechanism description has to be made (or to put it differently, the kind of entities included have to be determined). Alternatively, we could allow for a two level mechanism. Either way, a decision has to be made, and it is a *pragmatic* decision about the answer that we seek, when asking for an explanation of some phenomenon. Adding top-down conditions will not help, because manipulating the system's activity (e.g. from resting to running, calm to panic, etc.) will yield both the heart's pumping faster and the pacemaker cells generating more electric impulses per minute.

Second, the MM criteria are not specific enough to determine the boundaries of the mechanism. Thus, for example, if the explanandum phenomenon is delivering 4.5 liters of blood per minute to the organs in a human being, rather than 5 l/min then the boundary of the mechanism is not clear-cut by any means. We could understand the mechanism as being comprised of the heart, the arteries, and veins and explain the phenomenon by alluding to the heart rate. Alternatively, we could include the complete regulatory system of the heart, i.e. baroreceptors, the vasomotor regulatory center in the hypothalamus, etc.. Wiggling the baroreceptor (bottom-up) would change the heartbeat and the blood delivery and running (top-down) would change the state of the baroreceptor. Therefore, the baroreceptor fulfills the MM criteria. As we can see delimiting the time and precision boundaries of a mechanism involves an additional *pragmatic* decision. Omitting or including (earlier) causes, most virulent in etiological mechanisms, where there is no natural starting point, but also in constitutive ones (like the cyclical mechanism for blood delivery), is not determined by the explanandum phenomenon and MM.

Third, specification of the *degree of precision* for the description of the mechanism is also not settled by MM. A mechanism can explain how something works in broad strokes, or it can be very precise. There are different virtues and drawbacks in different circumstances. If, for example, we want to explain how the mechanism for action potentials is the *same* in all living organisms, we would have to tell the story in quite general terms, invoking comparatively vague causal interactions and entities (such as sodium and potassium channels, the membrane, ions, depolarization, hyperpolarization, etc., as we find it explained in textbooks). However, sometimes we are more interested in telling how action potentials behave in a very precise manner (e.g. in human Purkinje cells), thus including the specifics of the ion channel types' distributions and their behavior in an action potential. For example the sodium voltage-gated channels Nav1.1, Nav1.6, Nabeta1 and Nabeta2 in human Purkinje cells (Schaller and Caldwell 2003, 2), an overview of this diversity can be found in the 800 page book by Bertil Hille (2001). There is often a trade-off between the precision of the description of the mechanism and the scope of the phenomenon that can be explained by a mechanism, but this is not always the case. In stark contrast to some parts of physics and chemistry, in biology generality usually requires omitting the details produced by the ubiquitous diversity of living things. Thus, while the exact description of a single chemical reaction, (e.g. $2H^{+}+O^{2-} \rightarrow H_2O$) will not have to be less precise in order to be applicable to myriads of other instances, biological explanations often have decrease in precision in order to encompass more instances, a fortiori more instances of different taxa. Mechanism types whose scope is larger than one organism at one time will mostly describe mechanism tokens that have a family resemblance. They are similar in the aspects described, but dissimilar in others. There is, however, no objective way of deciding where one family of mechanisms stops and another begins. Moreover, in even only slightly different explanatory contexts the same organisms could be tokens of different types.

Finally, failing on the very aim it was constructed to accomplish, the distinction between a mechanism and its background conditions cannot always be settled on the grounds of the mutual manipulability criteria. Even in Craver's own example, the case of brain areas performing a certain task, both using fMRIs (top-down) and manipulating with these areas (bottom-up) will not exclude the background condition of blood flow in these areas from the mechanism. Craver writes that “mere correlates of task performance [blood flow] cannot be manipulated to change task performance” (Craver 2007, 159). Although he is right in stating that often manipulating background conditions “tend to be non-specific” (2007, 158), like in the case of the heart pumping blood and some performed task by the brain, this does not have to be the case. In an ideal intervention, we could throttle the blood flow in a very restricted area of the brain so that it would change task performance very subtly.[13] Fine-grained intervention on perceived background conditions are indistinguishable from the parts we want to include in a mechanism. Our decision to exclude blood flow from the mechanism stems rather from more general ideas about the kinds of causal interactions we want in our explanation. It is not the case that blood flow plays no causal role, but that typically it does not play an *interesting* role in the context of neurobiology. Neurobiological explanations typically omit the parts that concern the

[13] The idea of such an intervention is not to block the blood flow that normally surges *after* the neuronal activity but to cut off the oxygen and glucose supply ahead of time so that the neurons will not be able to perform properly. Arguing that the top-down condition will only change the surge after the activity, so MM excludes the blood flow *surge*, as the bottom-up intervention and the detected changes are distinct, is futile: of course no backward causation can take place, and a prolonged ideal intervention into the blood flow (with repeated task performance) will render it in accordance with MM. (Recent findings (Sirotin and Das 2009) even suggest that blood is sent in anticipation of neural events that never take place which could even suggest that a very rapid replenishment is crucial even for short periods of neuronal activity).

energy delivery (whether it is blood flow, mitochondria activity, etc.), as they are not salient.[14]

All these points show that a satisfactory solution for PMERP is only attainable by combining the manipulability criterion with several pragmatic decisions. We are now in position to put together the ontic and pragmatic parts of the relevance relation for a mechanism.

6. The pragmatic-ontic approach to mechanisms

According to van Fraassen, any why-question consists of a triple $Q = \{P_k, X, R\}$, where P_k is the *topic* of the question, $X = \{P_1, \ldots, P_k, \ldots\}$ is the *contrast class*, i.e., a set of statements that includes the topic plus a set of other propositions that are false, and R is a relevance relation. The answer to a question is a proposition A that is *relevant* to *Q.* This means that A bears the relation R to the couple $<P_k, X>$. Therefore, the question 'Why P_k?' presupposes that P_k is true and that all the other elements in X are false. Furthermore, question P_k is asking for a certain *kind* of answer, namely an answer that bears relation R (to the couple $< \mathrm{P_k}, \mathrm{X}>$) (cf. (Van Fraassen 1988))[15]. A well-known criticism of van Fraassen's account is that he does

[14] An interdisciplinary community that is interested exactly in these kinds of questions is however organized in the International Society for Cerebral Blood Flow & Metabolism.

[15] I believe that following the Van Frassen construction of pragmatics of explanation and adding the determinants to his relevance relation is the most straightforward method. However, it could be that an explication of the same points in a more elaborate framework of erotetic logic is possible. For example Hintikka & Halonen

not give any further specification of *R*, which turns it into an 'anything goes' position (Kitcher and Salmon 1987). This is, however, not implied (as Kitcher has already subsequently acknowledged (Kitcher 2001, 76, fn. 6). While van Fraassen does not say much about the relevance relation, my goal here is precisely to explicate the relevance relation for mechanistic questions. The necessary supplementations are determining factors for *R*. These determining factors are only applicable context-dependently, but are in no way random, loose, or arbitrary. First, there is the ontic, causal-manipulative part of the relevance condition. The terms that appear in a mechanism description are entities and their causal interactions. (This excludes for example statistical explanations as in population genetics or structural explanation as in unifying electricity and magnetism in Maxwell's equations). The entities and the causal interactions appearing in the answer have to change the explanandum behavior if manipulated (as defined in CR1, CR1a, and CR1b). This ensures that there is a causal efficacy of these entities and causal interactions for the explanandum. Second, there is a pragmatic part of the relevance condition. *Lexical* (or *level(s)*), *temporal*, and *precision constraints* must be added to the ontic-manipulative condition in order for the boundaries of the mechanism to be determined. All of these conditions should then select a *unique* set of entities and causal interactions that brings about the explanandum behavior. One arrives at a level of precision sufficient to meet the explanatory relevance problem only if one includes both the ontic-manipulative as well as the pragmatic conditions. As the

---

(2005) 'Interrogative-Nomological Method' would exchange contract classes with 'queried constants' (2005, 32-35) , and the determinants would be general properties of the background theory used. However as their method is based on laws and deduction, it does not fit the account developed here.

descriptions of mechanisms are answers to questions, pragmatic salience conditions are part of their determining factors. These are the conditions that distinguish *ideal explanatory texts*, which in principle would explain phenomena in a causal-mechanical way, from *descriptions of mechanisms* that are *partial explanatory information*.

We finally arrive at the following characterization of mechanistic explanation: Mechanistic explanations are answers to questions determined by the triple determinants of the question $Q = \{P_k, X, R\}$, and the quadruple determinants of the relevance relation $R = \{L, P, T, M\}$; where P_k stands for the topic (including the scope) of the question; X for the contrast class; R for the relevance relation; L for the level(s) or lexicon of the descriptive concepts;[16] P for the precision or the tendency of precision of the explanation (e.g. as detailed as possible, as general as possible); T for the time-frame and time resolution of the explanans (starting and end points, changes should be accounted for in a seconds-, milliseconds-, or nanoseconds-scale); and finally M, the mechanistic relation, which stays the same in all

[16] As many before me have realized, talk of 'levels' is notoriously problematic (see e.g. (Machamer and Sullivan 2003)), therefore I would suggest to call this determinant the *lexicon*, a set of (often interconnected) concepts used in (a) certain discipline(s) at a certain time to explain phenomena of a certain domain. However, I would like the determinant L not only to determine to what paradigm or general theory a certain explanation belongs, but a more narrow lexical sub-structure e.g. organs, cells, molecules , sodium channels, Kv1.1-Kv1.8, etc.

mechanistic explanation.[17] What makes these explanations explicitly mechanistic explanations is two-fold: first, the explanation has to be more *fine-grained* than the description of the explanandum, and second it has to involve *entities* and *direct causal interactions*, i.e. it has to include several entities that interact *directly* with each other, and if manipulated, change the explanandum phenomenon. These are the cornerstones of mechanistic explanation. If no entities and interactions are involved in the description (but for example only differential equations), then there is no mechanism.[18] If the interactions of the entities are not direct (in Woodward's sense of direct causal interactions), then we would claim that this is not a complete mechanism, but only a mechanism sketch (see (Darden 2005, 361); and if wiggling the entities (according to CR1, CR1a and CR1b) would not change the explanandum phenomenon, then we would claim that this entity is not relevant. As I will show below, although these determinants are not independent of each other, i.e. a decision on one determinant restricts choices on other determinants), there are still enough

[17] *M* is the same for all mechanistic explanations but can be varied if other kinds of explanation are sought. In [self-reference removed for blind review], we sketch a more general approach that includes different kinds of explanations.

[18] Of course, differential equations can sometimes be interpreted mechanistically, as for example in the case of the Hodgkin-Huxley equations, but only supplemented with such an interpretation, it amounts to a mechanistic explanation. (However, a causal interpretation of the Hodgkin-Huxley equations may not amount to a mechanistic one, see (Weber 2008), (Craver 2008) and (Bogen 2008) for a debate on this very topic).

degrees of freedom to each of them, so that it is not the case that deciding on one *implies* a decision on another.[19]

[19] A recently published account of functional explanation in context (Couch 2009) bears interesting similarities to the approach presented here. Comparing the two, however, reveals some problems with Couch's interesting construction. His account of functional explanation attempts to combine a pragmatic approach to the description of functional kinds with a variant of Mackie's *inus* condition for the realizing structure of that function. However, Couch does not supply a tool for making the explanandum function F precise. The *scope* and *contrast class* of the explananda deliver exactly this. Moreover, his use of Mackie's *inus* conditions imports a characteristic of this account that is in conflict with a pragmatic approach. *Inus* conditions do not distinguish between background conditions and active or salient parts. (As the generic example goes, the short circuit, oxygen, and the lack of a sprinkler are all factors in one *inus* condition for burning the house down). Thus, being 'context sensitive', as Couch suggests, in the definition of the explanandum kind F does not suffice in order to exclude any background from the explanation. Accordingly, in his example, the 'capacity of the human structure S for 'light reception' F', he claims that the *inus* set is complete with the receptor and the visual pigment, and "[t]his is not to deny that the other components may be relevant to certain aspects of the functioning of the eye; but in giving the explanation we should see that the other components can be ignored" (2009, 256). His example of an irrelevant component, the iris, seems convincing as it is indeed not directly relevant, but take the clarity of the lens: a severe cataract can lead to the loss of the capacity of light reception, so isn't an at least minimally transparent lens a necessary condition? Moreover, dead receptor cells do not have this capacity either, so blood,

7. Example: The Mechanism of Transcription

Taking into account the pragmatic concerns discussed in the previous sections, it is obvious that calling a mechanism the “mechanism of transcription” is so vague that it is not a real candidate for a mechanistic explanation. However, before turning to further specifications, it is worth mentioning that it is possible to call a similar mechanism the ‘mechanism for RNA synthesis’. *Nota bene*, both names are extremely vague and could also mean different processes (e.g. RNA synthesis *in vitro* vs. *in vivo* DNA transcription), however, the interesting point here is that it is possible to convert some constitutional mechanistic explanations into etiological explanations and vice versa: RNA synthesis is about the mechanism leading up to an RNA molecule and transcription is about the mechanism underlying this phenomenon. Although this is not universally the case, at least if we do not want to have very contrived formulations of the explanandum, it does point to a first decision about the explanandum and therefore the explanation sought: constitutional vs. etiological explanation. Although this decision has the consequence that the further determinants of the relevance relation are to be formulated differently and therefore the criteria for inclusion will be expressed differently, in some cases, the entities and causal interactions that will appear in the explanation could be the same for both constitutive and etiological explanations.

The second decision to be made is the scope of the explanandum (or of the topic P_k). Do we want to take into account all living cells, just eukaryotes, just those of the black rat (*rattus rattus*), etc.? If we do want the scope to be as general as

---

oxygen, and nutrient supply to the eye is also necessary, etc.. Only pragmatic constraints on the explanans sought, like *Lexicon, Time,* and *Precision*, could exclude these conditions. Mackie’s *ideal* account cannot serve to resolve this issue.

possible, this will impose certain restrictions about the possible details that could be involved (as discussed in the third point of Section Five). This is an important point as we can now see that to some extent the scope of P_k is already restricting the choices for other determinants, *precision* and *time resolution*. If we want to tell a story about all cells, we have to stay coarse-grained on the kind of entities, causal interactions and time resolution we seek. As I mentioned above, the *precision* and *time constraints* are not *fixed* by the choice of scope, they are only restricted by it. Two examples of explanation of transcription in different standard textbooks of biology can illustrate this. For example, here is the most concise one in (Lodish 2000, glossary) “Process whereby one strand of a DNA molecule is used as a template for synthesis of a complementary RNA by RNA polymerase”. Here is one with a little more detail:

> Transcription begins with the opening and unwinding of a small portion of the DNA double helix to expose the bases on each DNA strand. One of the two strands of the DNA double helix then acts as a template for the synthesis of an RNA molecule. As in DNA replication, the nucleotide sequence of the RNA chain is determined by the complementary base-pairing between incoming nucleotides and the DNA template. When a good match is made, the incoming ribonucleotide is covalently linked to the growing RNA chain in an enzymatically catalyzed reaction.[…]
>
> The enzymes that perform transcription are called RNA polymerases. Like the DNA polymerase that catalyzes DNA replication, RNA polymerases catalyze the formation of the phosphodiester bonds that link the nucleotides together to form a linear chain. The RNA polymerase moves stepwise along the DNA, unwinding the DNA helix just ahead of the active site for polymerization to expose a new region of the template strand for complementary base-pairing. In this way, the

growing RNA chain is extended by one nucleotide at a time in the 5′-to-3′ direction. The substrates are nucleoside triphosphates (ATP, CTP, UTP, and GTP); (Alberts et al. 1994, 241ff.).

As we can see, even if we want to describe the mechanism responsible for most of transcription (excluding e.g. reverse transcription), the picture can be more or less detailed. Additionally, the time frame we are interested in could be all of the process or just the beginning, i.e. the initiation process or another part of the mechanism.

Decision on the *Contrast class* focuses the explanation on different aspects of the process. We could ask when the transcription is successful, in the sense of accurate rather than inaccurate nucleotide sequence (according to usual complementary base pairing), which leads to a more detailed description of the proofreading part of transcription (not mentioned in the basic textbook description). We could ask for success in the sense of a specific mRNA to be transcripted, which would lead to a more detailed picture of the initiation process including upstream DNA enhancers and promoters, and the RNA Polymerase II complex (extending the spatial focus of the general explanans).

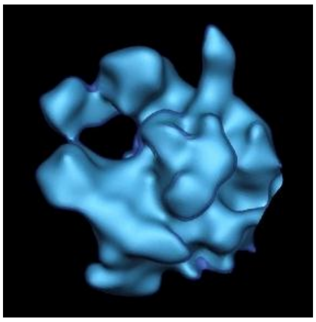

Figure 1: The RNA Polymerase II in yeast in a 16Å resolution (Darst et al. 1991, better quality image courtesy of Seth Darst)

To exemplify the distinction between the *precision* of the description and the included *lexicon* consider the resolution of a protein attained by crystallography (usually measured in Ångström). This kind of precision is not a matter of levels. While lower resolutions might suffice for certain explanations, (e.g. the assembly of sub-units, or the location of co-factors) others may require higher ones. As shown in figure 1, seeing the protein helps us understand the phenomenon as it allows us to make inferences about the working principles and interaction sites of the polymerase with the DNA strand (darker spots). Thus, the form of the protein is a relevant feature of the mechanism, too. It allows for the localization of different regions of the protein, which if intervened upon, would lead to different effects. Further experiments, like amino acid substitutions, exploit this kind of knowledge (see (Dieci et al. 1995) for this kind of research based partially on (Darst et al. 1991)).

This leads us to a final point: the *level(s)* or *lexicon* used in the explanation. We could decide to explain the phenomenon by alluding only to whole molecules, to structural features of molecules (like alpha helices and beta sheets), to sub-molecular parts (like amino acids), or even to single atoms and their causal interactions. This will mainly depend on the methods we use to elucidate the mechanism (see [self-reference removed for blind review]). However, one does not have to assume that the finer-grained entities are *mechanisms* that bring about the upper level entity. Continuing with the example above, the amino acid at a certain interaction site does not have to be part of a fully-fledged mechanistic explanation of the polymerase activity. It is sufficient that the amino acid residue in question is doing some causal work in the *overall* mechanism of transcription in order for it to be granted membership in the mechanism. As we can see, the resulting picture seems far more complex and messy than Craver's Matryoshka nesting dolls view, which proposes a more tidy 'levels of mechanism view' of mechanisms inside of mechanisms, inside of mechanisms (see (Craver 2007, ch. 5-7), and (Craver and Darden, 2001, 118). However, the biological literature suggests that biologists are quite opportunistic in their selection of entities (or lexica) in a mechanism. On the other hand, if one looks at *heuristics* in science (as some of Darden's work has shown), one can follow a string of questions that follow hierarchical patterns (see [removed for blind review] for further elaboration of this point).

To sum up, decisions on the five determinants of a mechanistic explanation (P_k, X, L, P, T) (presupposing the mechanistic relation M that makes the explanation mechanistic and R being determined by the quadruple) are made on *pragmatic* grounds, and only once they are fixed are the entities and 'activities' fixed as well. Thus, a mechanism should be understood as the set of the referents of the concepts used in an explanation. Mechanisms are not "acting entities" in any sense similar to a moving atom or an exploding building. To think that mechanisms are 'things' that have some kind of natural existence, akin to natural kinds, is a mistake. What exists, if at all, are the entities and their causal interactions.

References

Alberts, Bruce, et al. (1994), *Molecular Biology of the Cell, Third Edition*. New York: Garland.

Beatty, John (1997), "Why Do Biologists Argue Like They Do?" *Philosophy of Science* 64:S432-S443.

Bechtel, William, and Robert Richardson (1993), *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton, N.J.: Princeton University Press.

Bogen, Jim (2005), "Regularities and causality; generalizations and causal explanations", *Studies in History and Philosophy of Biological and Biomedical Sciences* 36:397-420.

Bromberger, Sylvain (1966), "Why Questions", in Robert Colodny (ed.), *Mind and Cosmos: Essay in Contemporary Science and Philosophy*, Pittsburgh: University of Pittsburgh Press, 86-111.

Couch, Mark (2009), "Functional Explanation in Context". *Philosophy of Science 76: 253-269.*

Craver, Carl and Lindley Darden (2001) "Discovering Mechanisms in Neurobiology: The Case of Spatial Memory", in Peter Machamer et al. (eds.) *Theory and Method in the Neurosciences. Pittsburgh University Press, 112-137.*

Craver, Carl (2007), *Explaning the Brain. Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Clarendon Press.

Darden, Lindley (2005), "Relations among fields: Mendelian, cytological and molecular mechanisms", *Studies in History and Philosophy of Biological and Biomedical Sciences* 36:349-371.

——— (2006), "Strategies for Discovering Mechanisms: Construction, Evaluation, Revision", in Lindley Darden (ed.), *Reasoning in Biological Discoveries. Essays on Mechanisms, Interfield Relations, and Anomaly Resolution*, CUP, 271-312.

Darst, Seth, et al. (1991), "Three-dimensional structure of yeast RNA polymerase II at 16 Å resolution ", *Cell* 66:121-128.

Delehanty, Megan (2005), "Emergent properties and the context objection to reduction", *Biology and Philosophy* 20:715-734.

Dieci, Giorgio, et al. (1995), "A universally conserved region of the largest subunit participates in the active site of RNA polymerase III". EMBO Journal 14:3766-3776.

Fehr, Carla (2004), "Feminism and Science: Mechanism Without Reductionism", *NWSA Journal* 16:136-156.

Glennan, Stuart (1996), "Mechanisms and The Nature of Causation", *Erkenntnis* 44:49-71.

Halonen, Ilpo, and Jaako Hintikka (2005), "Towards a Theory of Process of Explanation", *Synthese* 143:5-61.

Hille, Bertil (2001), *Ion Channels of Excitable Membranes. Third Edition*. Sunderland, Mass.: Sinauer.

Hitchcock, Christopher (1995), "Salmon on Explanatory Relevance - Discussion", *Philosophy of Science* 62:304-320.

——— (1996), "The Role of Contrast in Causal and Explanatory Claims", *Synthese* 107:395-419.

Kitcher, Philip (2001), *Science, Truth, and Democracy*. OUP.

Kitcher, Philip, and Wesley C. Salmon (1987), "Van Fraassen on Explanation", *The Journal of Philosophy* LXXXIV:315-330.

Kyburg, Henry (1965), "Comment (Salmon's Paper)", *Philosophy of Science* 32:147-151.

Lodish, Harvey, et al., (2000), *Molecular Cell Biology*.[internet] 4th ed. Basingstoke: Freeman. Available from ncbi.nlm.nih.gov/bookshelf/br.fcgi?book=mcb .

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 Craver (2000), "Thinking About Mechanisms", *Philosophy of Science* 67:1-25.

Machamer, Peter, and Jacqueline Sullivan (2003), "Leveling Reductionism", unpublished manuscript, http://philsci-archive.pitt.edu/archive/00000386/.

Mayntz, Renate (2004), "Mechanisms in the Analysis of Social Macro-Phenomena", *Philosophy of the Social Sciences* 34 :237-259.

Piccinini, Gualtiero (2007), "Computing Mechanisms", *Philosophy of Science* 74:501-526.

Psillos, Stathis (2004), "A Glimpse of the Secret Connexion:Harmonizing Mechanisms with Counterfactuals", *Perspectives on Science* 12:288-319.

Railton, Peter (1978), "Deductive-Nomological Model of Probabilistic Explanation", *Philosophy of Science* 45:206-226.

——— (1980), *Explaining Explanation: A Realist Account of Scientific Explanation and Understanding*. Dissertation. Princeton University.

Salmon, Wesley (1971), "Statistical Explanation", in Wesley Salmon (ed.), *Statistical Explanation and Statistical Relevance*, Pittsburgh: Pittsburgh University Press, 29-87.

——— (1984), *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.

——— (1997), "Causality and Explanation: A Reply to Two Critiques", *Philosophy of Science* 64:461-477.

Schaller, Kristin, and John Caldwell (2003), "Expression and distribution of voltage-gated sodium channels in the cerebellum." *Cerebellum* 2:2-9.

Sirotin, Yevgeniy and Aniruddha Das (2009), "Anticipatory haemodynamic signals in sensory cortex not predicted by local neuronal activity", *Nature* 457:475-479.

Thagard, Paul (2003), "Pathways to Biomedical Discovery", *Philosophy of Science* 70:235-254.

——— (2008), "Mental Illness from the Perspective of Theoretical Neuroscience", *Perspectives in Biology and Medicine* 51:335-352.

Van Fraassen, Bas (1988), "The Pragmatic Theory of Explanation", in Joseph Pitt (ed.), *Theories of Explanation*, OUP, 51-74.

Weber, Marcel (2008), "Causes without Mechanisms: Experimental Regularities, Physical Laws, and Neuroscientific Explanation." *Philosophy of Science 75:995-1007*.

Wimsatt, William C. (1976), "Reductive Explanation: A Functional Account", in R. S. Cohen, et. al. (eds.), *PSA 1974*, Dordrecht/Boston: Reidel, 671-710.