

Mechanisms, laws and explanation

Nancy Cartwright, John Pemberton and Sarah Wieten¹

1. Introduction

Mechanisms – ‘... entities and activities organized in such a way that they are responsible for [a] phenomenon’ – are now widely taken in philosophy of science to provide one of modern science’s basic explanatory devices.² Among the many different kinds of things they explain are the regular(ish) behaviors recorded in *ceteris paribus* (*cp*) laws: like ‘Ceteris paribus, the arrival of neurotransmitter particles at the head of a neuron is followed by the release shortly after of neurotransmitter particles from the synaptic vesicles at the other end’; or the recurring elliptical orbits of the planets around the sun recorded in Kepler’s laws. In 1999 Nancy Cartwright introduced the idea of a ‘nomological machine’ characterized as a ‘fixed (enough) arrangement of components, or factors, with stable (enough) capacities that in the right sort of stable (enough) environment will, with repeated operation, gives rise to the kind of regular behavior that we represent in our scientific laws.’ (Cartwright, 1999, page 50)³ The description of the nomological machine explains the law, which holds ‘ceteris paribus’ – *relative to* the nomological machine and its proper operation. Much of the more recent literature on mechanistic explanation of *cp* laws can be seen to make the same supposition when it comes to explaining *cp* laws (e.g. Craver 2007,

¹ Nancy Cartwright’s work for this paper is based upon research supported by the National Science Foundation under grant no. 1632471 and the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement no. 667526 K4U), for which she is very grateful. It is acknowledged that the content of this work reflects only the authors’ views and that the ERC is not responsible for any use that may be made of the information it contains.

² This characterisation from Illari and Williamson, 2012 is labelled a “consensus concept” by them and by the *Stanford Encyclopedia of Philosophy* article ‘Mechanisms in Science’ (Craver and Tabery, 2017). See that article for an extended discussion. Our characterisation of mechanisms in this paper is in accord with that, with Cartwright’s work on mechanisms and *cp* laws, and with the characterisations of the so-called ‘new mechanists’ (including much of the work by newer ‘new mechanists’). As Craver and Tabery report: “Three characterizations are most commonly cited:

- MDC: ‘Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions’ ...
- Glennan: ‘A mechanism for a behavior is a complex system that produces that behavior by the interaction of a number of parts, where the interaction between parts can be characterized by direct, invariant, change-relating generalizations’ ...
- Bechtel and Abrahamsen: ‘A mechanism is a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism is responsible for one or more phenomena’ ...

Each of these characterizations contains four basic features: (1) a phenomenon, (2) parts, (3) causings, and (4) organization.”

³ We suppose, like many ‘new mechanists’ (see e.g. Glennan, 2017, especially chapters 4 and 5), that individual mechanisms can be grouped under type labels and that the regularities we record as *cp* laws may arise from either the repeat operation of an individual mechanism, and hence be conditioned on the operation of that specific individual (like *the planetary system*) or from the operation of a plurality of mechanisms of the same type. In this latter case they are conditioned on mechanisms of that type (e.g. *neurons*) and could be expected to break down in any specific case where the specific token of that type misfires. We shall elide this distinction between types and tokens here because our points hold equally for both.

Bechtel & Abrahamsen 2005). In this paper we focus on how a mechanism explains the behavior recorded in a cp law.⁴

We will address two closely inter-related questions:

1. An epistemological question: “What kind of explanation is involved?”
2. An ontological question: “What is going on *in the world*?”

Our use of the term ‘mechanism’ accords both with Cartwright’s original characterization and with that of the so-called ‘new mechanists’ (see e.g. Illari and Williamson 2012, Glennan 2017) not only in spirit but in the main details. But since we look at only a single function – giving rise to and explaining regularities – among the many attributed to mechanisms, we specialize the characterization to suit this function. As Cartwright’s original definition has it, a mechanism is a *nomological machine* if, when operating without interference, it gives rise to stable input-output relations (regular behavior) of the kind typically recorded in *ceteris paribus* causal laws.⁵ For our purposes here we characterize a *nomological machine mechanism for regular behavior RB* by a set *P* of parts displaying a specific set *Y* of features (including activities and interactions) in an arrangement *A* which gives rise to behaviour *RB* if operating repeatedly without interference: $(M = \langle P, Y, A, RB \rangle)$.

To address the epistemological question, we enter into the lively debate about covering law versus mechanist explanation. We argue, contrary to what many mechanists themselves maintain, that when it comes to explaining cp laws, standard mechanistic explanations are a subset of good old-fashioned covering-law explanations (though a special subset) and that this is an important contribution to their explanatory force.

Explanation is a linguistic enterprise. For the ontology, we need to know what is going on in the world. We shall argue that the arrangement of the parts in the mechanism supplies them with features they do not possess separately. In the cases we focus on where covering laws play an important role in mechanistic explanation, what is happening in the world when $M (= \langle P, Y, A, RB \rangle)$ explains regular behavior *RB* is this: *RB* is what it takes for some set of principles that govern features in *Y* of parts *P* in arrangement *A* to be instanced in the operations *M* undergoes in giving rise to *RB*. So, given those principles, $(P \& Y \& A)$ implies *RB*.⁶

We speak to the two questions in turn in the next two sections, addressing in each case other answers available in the literature before setting out our own position. We stress the interrelatedness of our two answers.

2. Explaining ‘explains’ – the epistemic question answered

⁴ Some remarks about terminology. We suppose that in the cases we focus on, mechanisms – things in the world – give rise to regular behaviors, which are also things in the world. *Accounts* of what these mechanisms are like and how they operate *explain* these behaviors, which can be described in cp laws. For ease of expression we shall sometimes talk, as is not unusual in discussions in philosophy of science, of mechanisms or mechanistic explanations explaining the cp laws that describe regular behaviors.

⁵ We shall confine our attention to cp *causal* laws although mechanisms can give rise to behaviors described in non-causal laws as well. There is a vast literature on cp laws and on the factors nodded to in the cp clause. We focus here on the need to inter alia condition them on the proper operation of the nomological machine / mechanism that gives rise to them. The important topic of cp laws is addressed elsewhere by the authors – see e.g. Cartwright 2002, Pemberton & Cartwright 2014 – it is not explored further here. Also, we shall use ‘laws’ and ‘principles’ interchangeably, depending on what is common usage for the ones under discussion; and, as noted already, where there is no danger of confusion, we may not always distinguish cp laws from the behaviors they describe.

⁶ Thanks to an anonymous referee for urging us to make this implication explicit.

It might seem⁷ that there is an obvious answer to the question, ‘What sort explanation is involved in mechanistic explanation of a regular behavior?’ The answer is that mechanistic explanations explain by laying out the structure and operation of the mechanism that gives rise to that behavior, i.e. showing its parts, arrangements, interactions, activities and the like. But a vast array of mechanistic explanations across the natural and social sciences offer more by way of explanation. They make clear *why* that structure operating in that way gives rise to that behavior. They do so, we maintain, by invoking – sometimes explicitly, but often implicitly – familiar laws or principles under which the parts act. Our thesis is, then, that when it comes to explaining cp laws, a great many mechanistic explanations, including paradigmatic ones discussed by mechanist philosophers, are a subset of covering-law (CL) explanations and part of their explanatory force is due to that. They are, though, a special subset of CL explanations, where the specific facts invoked are facts about the parts, features, interactions and arrangements of the mechanism that gives rise to them, which gives them further explanatory ‘depth’, as demanded in point 2 below.

Though we suppose that the covering-law model of explanation is familiar to all readers, to ensure a common understanding we quote from CG Hempel’s ‘Two models of scientific explanation’ which lays out just what we mean by ‘covering-law explanation’ here:⁸

This explanatory account may be regarded as an argument to the effect that the event to be explained (let me call it the explanandum-event) was to be expected by reason of certain explanatory facts. These may be divided into two groups: (i) particular facts and (ii) uniformities expressed by general laws. ... If we imagine these various presuppositions explicitly spelled out, the idea suggests itself of construing the explanation as a deductive argument of this form:

$$\begin{array}{l} C_1, C_2, \dots, C_k \\ \hline L_1, L_2, \dots, L_r \\ E \end{array}$$

Here, C_1, C_2, \dots, C_k are statements describing the particular facts invoked; L_1, L_2, \dots, L_r are general laws: jointly, these statements will be said to form the explanans. The conclusion E is a statement describing the explanandum event; let me call it the explanandum-statement... (Hempel, 2002, pages 46-47).

Many mechanists see mechanistic explanation as very different from covering-law explanation. Stuart Glennan states that, “New Mechanists are of one voice in seeing mechanistic explanation as an alternative to covering law conceptions of explanation.” (Glennan, 2017, page 221.) Antti Revonsuo claims, “Explanation in basic neuroscience is a prime example of causal-mechanical explanation *rather than* explanation in terms of universal laws and principles.” (Revonsuo 2001, page 47. Emphasis added.) Or consider Craver and Tabery, who title the section on explanation in their *Stanford Encyclopedia* article ‘Mechanisms in Science: From Formal Analyses to Material Structures’, where the formal analysis in question is the covering-law account: “According to [the covering-law model], explanations are arguments showing that the event to be explained ... was to have been expected on the basis of laws of nature and the antecedent and boundary conditions ... Mechanists, *in contrast*, insist explanation is a matter of elucidating the causal structures that produce, underlie, or maintain the phenomenon of interest.” (Craver and Tabery, Section 3 (first para). Emphasis added.) We take it that in this remark “the event to be explained” is meant to include the events of our concern here – regular behaviors, and we agree that mechanistic explanations in the form of nomological machine models ‘elucidate the causal structures that produce, underlie, or maintain’ these regular behaviors – that was just the point of introducing them in the first place decades ago. But we disagree that this is at odds with the models showing “that the event to be explained ... was to have been expected on the basis of laws of nature and the antecedent and boundary conditions.”

⁷ As an anonymous referee suggests.

⁸ We also can include IS – inductive statistical – as well, as we note below. But it seems excessive to lay out that familiar definition too.

Craver and Tabery note four concerns expressed by mechanists about the covering-law model of explanation (Section 3.1):

1. its inability to deal with causal /etiological explanation;
2. its inability to distinguish re-descriptions of the phenomenon in general terms from explanations that reveal the mechanism that produces it;
3. its possible lack of depth (subsuming a phenomenon under any true law will count as a complete explanation so that the level of detail may be insufficient for satisfactory explanation);
4. its requirement for laws, which may often be unavailable in the biological and special sciences.

Practicing social scientists also sometimes contrast covering-law and mechanistic explanation. For instance, in their classic text *Case Studies and Theory Development in the Social Sciences*⁹, Alexander George and Andrew Bennett briefly review the standard philosophy of science literature and argue that mechanistic explanation can solve two problems faced by the deductive-nomological (D-N) account, which along with the Inductive-Statistical (I-S) account is the standard formulation of covering-law explanation. The first is the problem of distinguishing “between causal and spurious regularities” (page 132). This problem is akin to 1. above. The “second problem with the D-N model is that its predictions must be rendered with perfect certainty”, a problem which, they argue (following Wesley Salmon), the I-S version does not successfully solve. This adds another to the list of concerns about covering-law vis-à-vis mechanistic explanation:

5. For covering-law explanation, outcomes are supposed to be fixed by the particular facts and laws invoked.

We do not see such a contrast between mechanistic and covering-law explanation. We urge rather that a large array of the best mechanistic explanations for regular behaviors work in part *because* they fit this covering-law model. Standard mechanistic explanations of cp laws are not separate from covering-law explanations but are a subset of them, and a good number of exemplary covering-law explanations, including many that were offered from the very beginning when the covering-law (CL) account first came into fashion, are equally exemplary mechanistic ones. The explanation of Kepler’s first law, detailed below, which couples Newton’s laws with a description of the parts and interactions of the solar system, is a case in point.

For complete clarity let us note that our claim does not imply that everything that fits the CL model is a mechanistic explanation, nor that everything that fits the CL model is a good explanation, nor that every good mechanistic explanation of everything that might be explained by a mechanism is a CL explanation.¹⁰ Rather, to repeat, we claim that many good mechanistic explanations – including standard exemplars in the mechanism literature – of the behaviors described in cp laws fit the CL model and, moreover, that this is responsible for part of their explanatory force. The root idea behind the CL model is that the explanans explains the explanandum because the explanandum *is to be expected* given the explanans. Recall that we characterise a mechanism M by a set P of parts displaying a set of features γ in an arrangement A. We claim then that generally the repeated operation of mechanism M explains the regular behavior RB in part because RB is what is to be expected given that a set of principles, λ , that govern features in γ all obtain at once in the arrangement A of the parts P. The principles, λ , fill in the second of Hempel’s categories and the description of the parts P, the significant features γ (including interactions) and the arrangement A fill in Hempel’s first category “the particular facts invoked”. It is the special nature of “the particular facts invoked” – facts about the underlying causal structure that gives rise to RB – that singles out this special subclass of CL explanations as mechanistic. Like many of Hempel’s examples,¹¹ these explanations are both

⁹ George and Bennett 2004. See also Beach and Pedersen 2013, chapters 3 and 4.

¹⁰ We do not posit a single universal criterion of explanatory correctness and are thus amongst those that James Woodward characterises as moving away from a strict universalist position (Woodward, 2014, Section 7, para 1).

¹¹ For instance, explanation of Kepler’s laws by the structure of the planetary system and the laws governing it.

CL and mechanistic at once: they show *what* the underlying structure that gives rise to the behavior *is* by being mechanistic and they show *why* that behavior is to be expected by being CL.

For the sake of discussion, we stay neutral about how to interpret ‘law’ in the CL account and about how law-like the laws involved are. These may be exact laws (if there are any such), they may be pragmatic laws as urged by Sandra Mitchell (Mitchell 1997), they may be causal laws as discussed throughout philosophy of science nowadays, they may be tendency laws as discussed in Cartwright and other powers theorists (see e.g. Cartwright 1999), or they may be cp laws.¹²

Perhaps the feeling of contrast results from focusing on a simple subspecies of covering-law explanations that do not, at least on the face of them, invoke mechanisms. For instance: “Why does this neuron transmit messages?” Because: “CP, all neurons transmit messages.” Perhaps the role for the parts and their arrangements is not transparent in the description of covering-law explanations, since these may get lumped under the expression ‘antecedent and boundary conditions’. Or perhaps these boundary conditions are conceived of too simply. For instance: To explain Kepler’s 1st law (the planets travel in elliptical orbits with the sun as a focus), we use Newton’s $\mathbf{F} = m\dot{\mathbf{v}}$. The boundary conditions include the value of m and an initial value of \mathbf{v} . Given these we can solve the differential equation to get the elliptical orbit once we have written down a specific function for \mathbf{F} . That does not look like a mechanical explanation. But of course, far more is necessary. We have to fill in the specific functional form of \mathbf{F} . For that we need the structure of the mechanism.¹³ It is common in presenting this Newtonian explanation to use a diagram like Figure 2, picturing a simple nomological machine with just two parts, a large object and a small object, arranged some distance apart. Their relevant features are the masses of the two objects, M and m , their separation \mathbf{r} , and the relative velocity of the small mass relative to the larger, which has a component $\dot{\mathbf{r}}$ along \mathbf{r} and $\dot{\theta}$ perpendicular to \mathbf{r} . Because of the features of the objects and their arrangement, the larger pulls on the smaller with a force GMm/r^2 . Now we can construct a proper, filled-in differential equation. If we start our explanation with that filled-in equation, the role of the parts and their features, arrangements, and activities will not be apparent.

¹² Often, as an anonymous referee remarks, M is or seems to be in place and operating and the appropriate input occurs but the output does not, so B (‘input causes output’) does not occur. That can happen even when covering laws are involved in the appropriate way for a number of different reasons: because M does not really obtain (some parts or features or arrangements are flawed), something interferes with M ’s operation, the laws are only tendency laws or they are stochastic or they are even more permissive (e.g. they dictate a range of outcomes but don’t fix a probability over them). This last is allowed on Mitchell’s pragmatic view. It was discussed by GEM Anscombe in her Cambridge Inaugural Lecture (Anscombe 1971) where she argued that laws may describe a cause that is ‘enough’ for the effect to occur but is not logically sufficient for the effect. It is also defended by Pemberton & Cartwright 2014.

¹³ You could just treat the specific functional form as yet another ‘boundary’ condition. In this case your CL explanation is not properly mechanistic. But the fuller explanation we cite here – which is the standard one – is mechanistic: it takes the parts and arrangements and interactions as the boundary condition and uses the *law of gravity* to assign \mathbf{F} .




Figure 2: Elliptic orbit of small mass around large mass

So, Kepler's laws are deduced from (i) the initial conditions which in this case comprise a description of the components of the mechanism – a large mass M (the Sun) and a small mass m (the planet) and their arrangement – separated by r and far enough from other sources of force to make any but the gravitational attraction between them negligible; and their interactions or activities – mutual gravitational attraction; and (ii) Newton's laws, including the general principle that $\mathbf{F} = m\mathbf{\dot{v}}$ and the law of gravity that an object of mass m located r from another mass M experiences a force GMm/r^2 . The Phillip's curve, which pictures a trade-off between inflation and unemployment, is deduced by Chicago School economists in a model from (i) the ascription of beliefs to business agents that inflation-related price rises indicate increased real demand for their product and (ii) the general principles that agents act to maximize their expected utility, along with the bridge principle that an entrepreneur's utility in the setting modeled is constituted by the firm's profits.¹⁴

Objections 1,2, and 3 adumbrated by Craver and Tabery are thus not relevant to our thesis. It is not the case that CL explanations cannot do the jobs Craver and Tabery cite. But they don't do these jobs just by virtue of being CL alone. Like mechanistic models of explanation, CL models can do a variety of different things. What exactly a specific CL explanation does beyond what they all do – show what is to be expected given the laws of nature – depends on the features of the specific facts and principles invoked in it. We do not claim that any derivation satisfying the general covering-law demands can do the jobs Craver and Tabery call for. Rather, we claim, if various standard mechanistic explanations can do these jobs, as they argue, then so can (many) covering-law explanations since these standard mechanistic explanations of cp laws are a subspecies of covering-law explanations. Nor does objection 5 bear on our claims for we do not suppose that covering laws are all either 'deterministic' or statistical. Rather many of the central covering laws used in mechanistic explanations are 'tendency laws' (like the law of gravity and Coulomb's law) that tell what a cause contributes to the effect, not what overall effect actually happens, as in the law of gravity,

¹⁴ For yet a different kind of example consider CG Hempel's own reconstruction of John Dewey's explanation of some peculiar behavior he observed in soap bubbles in which the particular facts that enter into the explanans describe the parts involved and their arrangement: "the tumblers had been immersed in soap suds of a temperature considerably higher than that of the surrounding air; they were put, upside down, on a plate on which a puddle of soapy water had formed that provided a connecting soap film, and so on" (Hempel 1970, page 336). Note also Hempel's remark in the famous short text *Philosophy of Natural Science*: "As this use [to explain Kepler's and Galileo's law] of Newton's laws illustrates, empirical laws are often explained by means of theoretical principles that refer to structures and processes underlying the uniformities in question" (Hempel 1966, page 51. Emphasis added).

Coulomb’s law and the law describing the drag of the air in the Millikan experiment described above.¹⁵ This leaves objection 4.

In regard to objection 4, we agree that many advances in biology have little to do with the discovery of new general laws and almost everything to do with uncovering the structure of systems. But that has no bearing on whether general laws play a central role in the models biologists construct of how those systems do what they do. However the model is presented, with diagrams and figures, or with equations, or narratives, or whatever, why should we believe that structures that match the model can do what they are supposed to? The answer, we propose, is very often that these effects are just what is to be expected given the features of the parts in that arrangement and the covering laws in which these features figure. *The effects are just what is to be expected because that is what must happen if all those features act as they should under the general laws that govern them.* Each of the features – like the mass of the sun and the mass of a planet –acts in accord with the general laws that govern it. Their joint actions, in accord with all these general laws at once, explain why the system does just what it does.

It may be thought, “Yes, but physics examples are the easy ones. It is no surprise that physics, with its rich tool-kit of general principles, uses covering laws in its mechanistic explanations. What about elsewhere?” We maintain that physics is not special. Examples of mechanistic covering-law explanations of cp laws in the socio-economic sciences, described as such, can be found in Cartwright 1995. These include a money multiplier and a debt-generating mechanism. Here we consider one of the mechanists’ own favorite examples from biology: signal transmission in the neuron, where the regularly (though *far* from exceptionlessly) occurring behavior to be explained is that the arrival of neurotransmitter particles at the head of the neuron leads to the release of neurotransmitter particles from the synaptic vesicles at the other end of the neuron shortly after. Figure 3 illustrates a few of the parts and stages involved in such transmission.




Figure 3: Some steps in signal transmission in a neuron

Beginning students are typically told, “Signals *within* neurons are transmitted electrically, however signals *between* neurons are transmitted chemically across the synapse.” (Flannigan et al. 2015, page 117.) We take it that this means that these signals are transmitted in accord with well-known laws of electricity (such

¹⁵ For more on tendency laws see, for example, Mill 1967; or Cartwright 1989, chapter 4.

as Coulomb's law) and of chemistry. Here is the basic mechanistic explanation for the electromagnetic transmission within the neuron.

Positively-charged sodium and potassium ions sit both outside and inside the neuron. In the rest state there is somewhat more positive charge on the outside than inside, so the voltage measured from inside is slightly negative. Arriving neurotransmitter particles dock with receptors opening sodium and potassium gates in the neuron wall, allowing ions to enter the neuron, increasing the positive charge in the cell and the local voltage across the cell wall. If the voltage passes a threshold, that stimulates an adjacent sodium gate to open; sodium ions outside flood through the gate due to the electro-chemical gradient. The change of charge distribution increases the local potential difference across the wall, opening adjacent gates. In the meantime, the open sodium gates close quickly stopping the rise in voltage. The stimulus also opens a slower-to-open potassium gate and positively-charged potassium ions flow out because the membrane is now more negatively charged on the outside than the inside, so the voltage drops. The potassium gate then closes. Sodium and potassium pumps then restore the rest state so that the process can be repeated.

There is of course more to be said; for instance, concerning the diffusion forces affecting the flow of the sodium and potassium ions. But these too behave as they should, according to standard diffusion equations. Or, how do the gates open and close? A helical protein string embedded in a pore in the neuron wall features an uneven charge distribution and is contorted by this potential difference in a way that causes the opening of the passage. Again, this is what is to be expected given the basic laws of electromagnetics – which is not surprising since signal transmission across the neuron is modelled *electrically*, as the textbook cited says. Beyond that, one might explain how the protein is structured that allows it to contort as it does. That is likely not to use electromagnetic principles. But for satisfactory explanation it should use general principles that hold not just in proteins in neuron gates but elsewhere as well.

An explanation like this is both mechanical *and* covering law. The use of covering laws is what makes it an *explanation* and not a mere description of what happens. We understand why the explanandum behavior occurs given the structure of the mechanism because that is the behavior that must occur if none of the laws cited, which we take to be true, are violated in this mechanism.

Besides objections 1-5 above, others have been raised as well. Kaplan and Craver (Kaplan & Craver 2011, page 607) criticize 'predictivism' and in the same breath the CL model. But we do not claim the symmetry of explanation and prediction and in particular not that everything (like the prediction of the storm by the fall in the barometer that Kaplan and Craver cite) that might fit a CL model is a satisfactory explanation. This 'symmetry' claim is back to front from our claim that many satisfactory explanations – especially many mechanistic ones—are CL explanations and moreover that part of their force as explanations derives from their use of covering laws to derive the explanandum phenomenon. Kaplan & Craver (page 608) also object that a law like Snell's that says 'Light does X' does not explain *why* light does X. But that too is irrelevant. The kind of CL explanations of why the neuron transmits an action potential do not just invoke a law that says that (cp) they do so.

Holly Andersen (Andersen 2011) takes issue with Bert Leuridan's (Leuridan 2010) claim that laws (he thinks in terms of Mitchell's 'pragmatic' laws) are ontologically necessary to mechanisms,¹⁶ which resonates with ours that the regular behavior that is explained by mechanisms generally depends on the operation of laws and that this dependence is in part what gives mechanisms their explanatory force. Andersen allows that some sciences may not use mechanisms as explanatory devices but may rely only on laws. But she seems to deny our claims when she says "Mechanisms... are an alternative to laws as an explanation of regularities," and further claims, "Within sciences like biology and neuroscience, however, there is a fairly

¹⁶ See also Craver & Kaiser 2013.

straightforward way in which law-based and mechanism accounts of explanation are in direct competition” (Andersen, 2011, page 329). But then it seems that by “law-based explanations” she has in mind laws that “are based on stable regularities” (page 326). This then seems like laws like Snell’s law that Kaplan & Craver discuss. In that case we agree. Explanations of regular behavior by laws that describe that behavior as law-like and not accidental do not do the same kind of explanatory job as mechanisms and the laws engaged in their production of regular behavior.

Again, that not every CL-form deduction is a good explanation and that not every CL-form deduction is a mechanistic explanation is tangential to whether standard mechanistic explanations of cp laws are CL and that being CL is part of how they go beyond showing *what* causal structure gives rise to the regular behavior to show *why* it does so.

Next we turn to three conventional objections to CL models that Craver raises: “(i) the problem of distinguishing laws of nature from accidents and other non-explanatory generalizations; (ii) the problem of providing an account of explanatory relevance; and (iii) the fact that one need not show that a phenomenon was to be expected in order to explain it”(Craver, 2007, page 35).

(i) is not an issue for whether mechanistic explanations are instances of CL explanation. There are a variety of ways on offer to select out what counts as a scientific principle, including ones we have already mentioned like Mitchell’s pragmatic laws account and the tendency-law account, together with the invariance account of Woodward that Craver endorses. As best we can reconstruct, Craver himself addresses the problem of explanatory relevance in (ii) by way of constitutive relevance (Craver, 2007, pages 139-160).¹⁷ We, similarly, can appeal to an ontological account of the mechanism (as set out in Section 3) as providing a basis for a solution to the problem of explanatory relevance. In relation to objection (iii), Jim Bogen argues (following GEM Anscombe) that “causality is one thing and regularity is another” (Bogen 2008, page 112). Craver notes that “We have no difficulty in imagining quite irregular mechanisms, such as the mechanism for neurotransmitter release, that works roughly ten percent of the time, or a rusty chainsaw that starts arbitrarily infrequently” (Craver & Kaiser 2013). In such cases, the phenomenon to be explained is not to be expected given the mechanism. There are at least two responses to (iii) that are readily available on our CL view. One is that the relevant CL principles may be indeterministic or permissive in various other ways (recall the discussion in footnote 12 of the variety of visions in contemporary philosophy of science of what scientific principles and laws might be like).¹⁸ The other is that the requisite features or the requisite arrangement that characterize the appropriately operating mechanism are not in fact in place in the cases when the explanandum behavior does not occur. For example, a cistern may not flush on some occasions when the handle is pushed if bolts attaching the lever arm to the lifting rod become too loose.

Last, Jim Bogen urges that what matters in mechanist explanation is the detailed descriptions of the activities, which will be cases of causings but very thickly described. He proposes that what matters in each instance is what these activities are and what they do in that instance and not what happens elsewhere on other occasions (e.g. in Bogen 2008). This seems to threaten the idea that the activities generally fall under more general causal principles. When we look at his cases, however, it seems as if the opposite is true, just as in our discussion of the neuron example. For instance, Bogen concludes his example of the exploration of fermentation in the latter half of the nineteenth century: “Fermentation turned out to be a physico-chemical process” (Bogen, 2008, page 121) – that is to say, as we understand him, that the process follows general physical and chemical principles. Perhaps it might not always seem that general causal principles are in play because the thick descriptions of the activities are not always those that appear in the relevant causal principles, since one and the same activity can be described in different ways and at different levels of abstraction. For instance, pressing down the handle on a toilet cistern raises the other end of the lever

¹⁷ Craver’s proposes a mutual manipulability criterion to resolve the important but vexed question of constitutive relevance – see also Baumgartner and Casini 2017, Baumgartner and Gebharter 2016, Couch 2011, Harinen 2018.

¹⁸ See footnote 12 for more on why ‘expected’ behaviour may fail.

which pulls up the lifting rod which in turn pulls open the valve – the operation of the cistern follows from general principles such as the law of the lever, but there is no general principle that pressing a handle opens a valve.

We do not want to be dictatorial about the term ‘explanation’ though. There may be cases in biology and elsewhere where features in a mechanism do not obey the general principles that they do outside it. There may also be cases where singular causings do not fall under wider principles or laws.¹⁹ Either of these last two could be what is meant by the claim that the behavior that the mechanism gives rise to is *emergent*. We do not want to deny that there may be emergent behavior in either of these senses. In cases like these, describing what is going on in the mechanism when the explanandum behavior occurs is certainly a contribution to knowledge; perhaps it is reasonably called ‘explanation’. What we want to stress is that, by far and away, most of the satisfying mechanistic explanations in both natural and social science are covering-law explanations: we use already familiar laws and principles to show that the behaviour is to be expected in settings with just this structure. It is true that some covering-law explanations are not mechanistic explanations. The Hodgkin–Huxley formalism often cited by mechanists or the differential equations describing transitions between neuron states in a Markovian scheme are good examples. But, to repeat, that some covering-law explanations are not mechanistic does not show that most mechanistic explanations are not covering-law.

Recall the job that adding covering laws to the description of a mechanism and its operation is supposed to do. The description of a mechanism and its operation shows what structure and activities it is that when repeated give rise to the regular behavior to be explained. But that by itself does not show why they do so. We observe that in a host of actual mechanistic explanations on offer in science, this job is done by implicitly or explicitly invoking laws and principles that are taken to hold generally for the features of the mechanism (parts, properties, interactions, activities). We have allowed that sometimes, as in cases one might label ‘emergent’, it is just the case that the interactions and activities cited *do* give rise to the behavior in question (which can be verified in a variety of ways) with no more to be said. We are also happy to allow that there may be other ways to show why just those interactions of parts with those properties in that arrangement give rise to that behavior, though we haven’t come up with any good suggestions ourselves. We only claim that a great many mechanistic explanations given in science do it the CL way (though perhaps without much note). Supposing, then, that we are right that the covering-law model is generally in play in good mechanistic explanations, that answers our *epistemic question*. Next we turn to our *ontological question*.

3. The ontological question answered

Our introductory remarks talked about M ‘giving rise to’ the behavior B that is described in the cp law. This is not the only terminology in use. For Peter Machamer, Lindley Darden, and Craver (*MDC*), M is “productive” of B (Machamer, P. et al. 2000, page 3). According to Stuart Glennan, M “produces” B (Glennan 2005, page S344). William Bechtel and Adele Abrahamsen say that the operation of M is “responsible for” B (Bechtel and Abrahamsen 2005, page 423). Craver and Tabery 2017 add “underlying” and “maintaining” (Section 2.1.1.). There are, of course, differing views of the relation between the mechanism and the causal regularity it gives rise to. We focus on three leading treatments available in the literature, each of which has problems.

Causes 1. We might think in terms of a word that often slips into discussion of this issue and that is suggested by the terminology of production: *causes*. This word takes on different guises in different circumstances. Here we think it is unhelpful. In what sense does the operation of a mechanism *cause* the

¹⁹ So: Event 1 causes event 2 but there is no description C of event 1 and E of event 2 such that ‘C’s cause E’s’ is a law or principle (of some sort or another). This is explicitly denied by Donald Davidson (Davidson 1963, 1967) but we do not want to rule it out.

regular behavior B it gives rise to? For concreteness, consider the regular behavior expressed in a cp law of the form 'Cp, Fs cause Gs'. How, for example, does *the operation of the toilet mechanism cause pressings of the lever to cause the flushings of the cistern?* Generally, causes proceed their effects. But the operation of the mechanism and the causal process it gives rise to are simultaneous. Most causal processes are continuous in time and thus have intermediate steps. Interrupting these is a conventional strategy for preventing an unwanted effect once its cause has occurred. How does that work when the cause is the operation of a mechanism and the effect is the causing of G by F? Also, there should generally be a flow of influence from cause to effect, so that perhaps we can mark the putative cause – in this case, that would be the nomological machine, say the toilet mechanism -- and find the mark later on the effect – in this case, the causing of the toilet to flush by pressing the lever. None of these conventional characteristics of a causal relation are easy to find here. So this does not seem a promising starting idea.

Causes 2. A second strategy that assumes the mechanism, let's call it M, plays a proper causal role is to claim that the cp law is under specified. A full specification puts the mechanism into the antecedent of the law itself: M and F cause G. This is the approach that Judea Pearl among others advocates in his work on causal Bayes nets (Pearl 2009). It has a number of problems. If it is the operation of the mechanism that M is supposed to represent, then F in the antecedent is redundant: if the cistern mechanism operates, the lever is pressed. If it is the parts and their arrangement that M represents, are we to think of the parts and the arrangement as a cause? This is what Pearl seems committed to since on his proposals M figures into the causal graph and into the causal equations in just the same way as F. Yet if M is a cause, we could expect it generally to have the kinds of characteristics of causes we just described. But again, it takes fancy footwork to maintain that these characteristics are there or explain away the need for them. The use of this proposal in Bayes nets faces the additional difficulty that the nodes in a Bayes net are supposed to be random variables. That means they have a range of allowed values with a probability distribution over them. But what are the allowed values in our toilet example? Any structure, dreamt or undreamt of, that has a lever? And where can the probabilities over these come from?

Constitution. This is what is advocated by Craver, according to whom mechanistic explanation "is *constitutive* (or componential)" (Craver 2007, page 8. Original italics). Craver uses the diagram in Figure 1 to represent the connection between the phenomenon and the mechanism, stating that:

S's ψ -ing is explained by the organization of entities $\{X_1, X_2, \dots, X_M\}$ and activities $\{\phi_1, \phi_2, \dots, \phi_M\}$ (Craver 2007, page 7).

Where S is the mechanism, ψ is the behavior of S 'as a whole', the X_i 's are the component entities of S, and ϕ_i is the activity of component i.

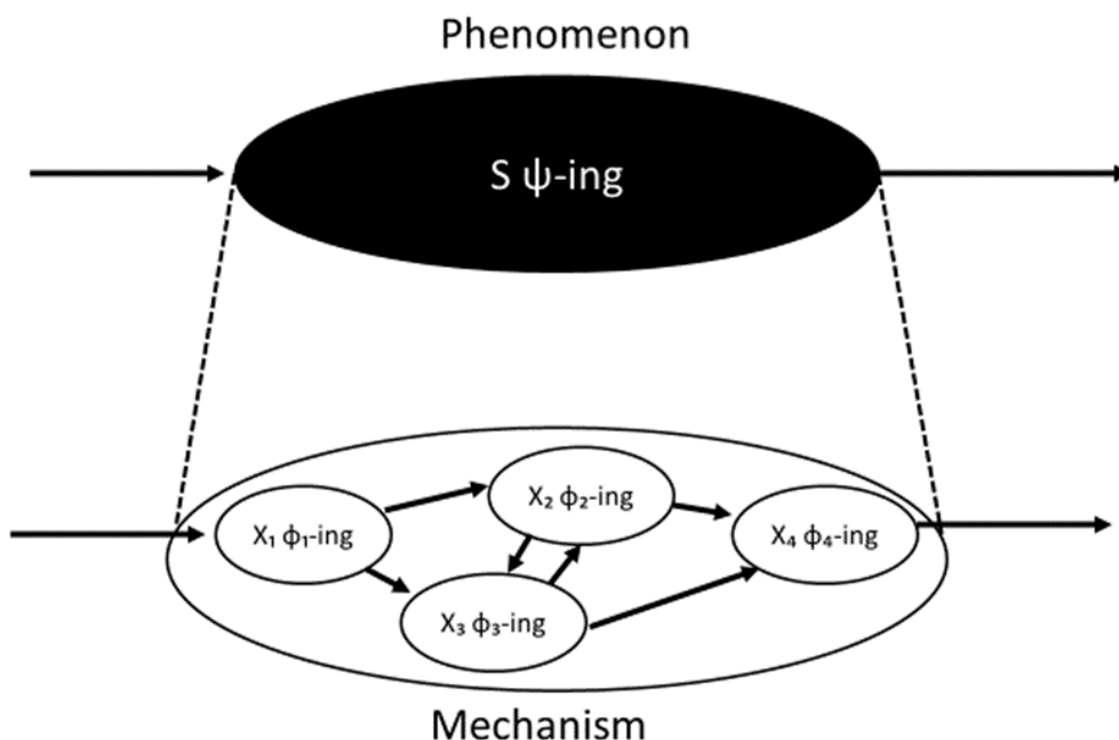


Figure 1: Craver's diagram of a phenomenon and its mechanism (Craver 2007, page 7, Figure 1.1)

This constitution account certainly avoids the difficulties facing causation and it seems to work for the kinds of cases that Craver focuses on, where the phenomenon is the ψ -ing of a system S , and the mechanism is the organized activities of the parts of S itself at each stage of S 's ψ -ing. The canonical example is the neuron transmitting a signal, which on Craver's account just is, or is constituted by, the organized components of the neuron and their activities, especially its membrane and gates and the potassium and sodium ions. In these cases, S and M refer to the same thing and, as Craver and Tabery put it in the *Stanford Encyclopedia*, "The phenomenon [to be explained] is the behavior of the mechanism as a whole." (Craver and Tabery, Section 2.1.) But what about other cases where the mechanism explains a cp law: 'CP, Fs cause Gs' where F and/or G may not be features of the mechanism but instead features of inputs and outputs to it. For instance, putting five quarters in the machine causes a can of coke to drop out (which is true of a vending machine but not of a parking meter). Here the mechanism is employed to explain how a single feature that is not a feature of the mechanism itself – F : the insertion of a quarter – causes a different feature – G : can of coke appearing in the reception bin – that may also not be a feature of the mechanism. Constitution makes sense in the neuron case. Typical synonyms for 'constitutes' are 'amounts to', 'adds up to', 'makes up', 'composes', and 'comprises'. It seems true that the neuron's parts doing what they do amounts to/adds up to/makes up/composes/comprises its transmission of a potential difference. It is difficult to see, though, how the causing of the appearance of the coke by insertion of the coins is made up of the operations of the vending machine.

But even if we accept that the activities of M constitute the causing of G s by F s, that would not solve our puzzle. We still need an account of *why*. Generally, when it is true that x constitutes y , there is a reason why it does so. When x constitutes y , that is not a brute, or isolated, fact. There are other facts without which x would not constitute y . The kind of reason can vary from case to case. What matters is that it is not arbitrary what constitutes what, or what kinds of things constitute what other kinds.

At the Board of Examiners meeting, the Chair takes role and announces, 'We constitute a quorum.' Why do we constitute a quorum? Because 'we' includes the Chair of the Board of Examiners, the Secretary, all three external examiners, and five internal members of the Board. That is what the University's *Learning and Teaching Handbook* says it takes to make a quorum. Later at the meeting you raise your hand after a proposal has been discussed. Raising your hand constitutes voting 'yes' to the proposal. It does so because that's the convention at the Examiners' meeting.

The reason of course need not be something written in a rule book. It can, for example, depend on the kind of thing that is to be constituted and what that thing is supposed to do. Why can't a heap of bricks constitute a fence? Because a fence is meant to enclose an area, and a heap of bricks does not do that. The parts in the arrangement pictured in the diagram in the design specifications constitute the toaster because that's what makes it up and allows it to do its job.

So, even if the relationship between behaviors in the mechanism and those described in the cp law is constitution, this leaves an unanswered ontological question, parallel to the one we ask. We ask, 'What is it for a mechanism's operating to give rise to B'. If the answer is that its operating constitutes B, what is the *reason* for that? *Why* do the joint activities of the parts of the mechanism in this arrangement constitute this behavior?

Our answer to our ontological question involves constitution, but not in the way that Craver pictures it. Ideas from Adolf Grunbaum (Grunbaum 1963) give a clue as to how. Newton's laws explain Kepler's because Kepler's laws are what Newton's *amount to* in the context of the planetary system. The behavior described in Kepler's laws *constitutes* the obtaining of Newton's laws given the features and arrangement of the planets and the sun. Travelling in the elliptical orbit prescribed by Kepler just is what it is for a planet to do what Newton's laws dictate in the presence of the sun. So, our suggested answer to the ontological question (which follows closely our answer to the epistemic question) is this:

Suppose that behavior RB occurs if mechanism M (characterized by parts P, arrangement A and features Υ) operates successfully repeatedly. $M = \langle P, A, \Upsilon \rangle$ gives rise to RB if, for some $\Upsilon' \subseteq \Upsilon$ and general principles $\lambda(\Upsilon')$ governing features in Υ' , all the principles in $\lambda(\Upsilon')$ are instantiated in RB's occurring.²⁰

John Pemberton's work stresses a point made widely by 'new mechanists': the importance of paying sufficient attention to arrangements (Pemberton 2011). Arrangements matter crucially here because they confine how general principles are instanced. For example, Towfic Shomar models an arrangement in which two charges attract each other yet the one moves away from the other, in part on account of that attraction; if it did not, Coulomb's law would be violated in that arrangement (Cartwright 1999, Figure 3.1a on page 60 and Figure 3.1b on page 61).

Arrangements play two roles in giving rise to regular behavior. First, they introduce new features that parts do not have by themselves. A strong branch, or a shovel, balanced over a rock or a log becomes a lever, which obeys the law of the lever, as levers do wheresoever a lever is found, whether with its end wedged under a wheel to heave a car out of the mud or functioning as a seesaw. Differing distributions of sodium ions inside or outside the neuron membrane exhibit differing voltages across the membrane. So, by virtue of the arrangement, new features obtain and new laws are called into play. Second, arrangements fix which activities happen when: which happen together and which after which. In the neuron, the gates nearest the incoming impulse open first, then the next ones along, so that the gates open sequentially along the length of the axon, which is typically long and thin. One could imagine a differently-shaped membrane with a more symmetric arrangement, which would give rise to very different behavior in the mechanism. Together these mean that the arrangements are crucial to what general laws are instantiated in the mechanism and

²⁰ That is, RB is an instance of all those principles holding, just as the earth's travelling in an elliptical orbit is jointly an instance of $F = ma$ and $F_G = GMm/r^2$.

what the behavior is when they are all instantiated in the same process. This is why arrangements play such a central role in mechanistic explanation. But they do so in part precisely because mechanist explanations rely on covering laws.

We suppose that the general laws relevant to the features of the mechanism determine the behavior of its parts within each arbitrarily short time period. These laws may be expressible by differential equations; or laws of heating, compressing (e.g., laws concerning coefficients of restitution), stretching, distorting, retarding (e.g. laws concerning friction), dissolving, diffusing, etc. Other laws may be expressible qualitatively (e.g., laws governing the cutting of a knife). When the mechanism operates normally, these laws obtaining simultaneously for all the parts together in their given arrangement determine the behavior of the parts of the mechanism in each arbitrarily short time period and hence the continuous behavior of the mechanism through time. Consider the case where input F causes output G, as in the quarters and the coke can. Together the laws obtaining for all the parts of the mechanism together are the reason that the initial state of the mechanism causes the final state and hence that the stimulus F that causes the starting state can truly be said to cause G later.

Consider neuron transmission again. Let $F(t)$ be the arrival at t of neurotransmitter particles at the head of the neuron; $G(t')$, the release of neurotransmitter particles from the synaptic vesicles at the end of the neuron. $F(t)$ causes sodium and potassium gates to open. This is part of the starting state of the neuron. Other important aspects of the starting state are that the voltage-gated channels are closed and there are more potassium ions inside the neuron than outside and conversely with sodium ions. The neuron then exhibits a continuous, orderly sequence of states over time, sparked by $F(t)$, crucial among them being ones which exhibit a potential difference above threshold, which cause the first sodium-gated channel to open, which causes later states in which others open, in turn producing a neuron state in which there is a large potential difference. So the action potential travels down the neuron's axon to the presynaptic terminal at the end. That in turn causes $G(t')$.

Although in Section 2 we focused on Coulomb's law, in this process we see activities in which a number of different well-established general laws are co-instantiated, for instance:

- a) A cloud of particles contained by a wall in which there is a gate which is open/closed) can enter/not enter the gate and cross the wall.
- b) A (net) force on a free-moving particle accelerates/moves it in the direction of that force.
- c) A distribution of charges gives rise, via Coulomb's law, to forces on local charged particles.
- d) A flexible object subject to differential forces on differing parts contorts.

Although these laws are familiar and unremarkable, they are central to the operation of the neuron, as to many other mechanisms.

We note that our answer to the ontological question is orthogonal to the issue of whether, and in what sense, mechanisms must be able to give rise to regular behavior (see e.g. Andersen 2012, Glennan 2010 and Krickel 2017) since our claims are about those mechanisms that do give rise to regular behavior. We allow that there may be, as Glennan (2010) claims, one-off or ephemeral mechanisms in areas such as history, for example. Although these are outside the scope of our study, we can note that in many cases of historical explanation, the behavior explained by the mechanism is nevertheless what it takes for all the principles governing the features of the mechanism to be instantiated together.

We note too that our answer is neutral concerning other ontological issues associated with mechanisms that have been the topic of extensive debate amongst 'new mechanists'. For example, many 'new mechanists' have offered accounts of the nature of causal relations which they take to underpin mechanisms – Woodward's manipulationist approach has proved popular (see e.g. Bogen 2005, Craver 2007 chapter 3, Craver and Bechtel 2007, Glennan 20011, Glennan 2017 chapter 6). Others have explored questions concerning the ontological nature of the mechanism, e.g. to what metaphysical category of being

does it belong? (see e.g. Kaiser and Krickel 2017, Krickel 2018). Here we remain agnostic as to whether the general principles instantiated in the mechanism which give rise to its behavior may be derived from change-relating generalizations, Humean regularities, powers, or other causal or non-causal principles and whether they may be associated with substances, events, occurrents, continuants or entities in other posited metaphysical categories.

4. Conclusion

When a mechanism explains a cp law, generally (a species of) covering-law explanation is in play. When a mechanism M gives rise to a regular behavior RB that is described in a cp law, RB is what it takes for some set of principles that govern the features of M's parts in their arrangement in M all to be instanced together. These answers to our epistemological and ontological questions, respectively, are closely interrelated.

References

- Andersen, H. K. (2011). Mechanisms, laws, and regularities. *Philosophy of Science*, Vol. 78, No. 2, pp. 325-331.
- Andersen, H. K. (2012). The case for regularity in mechanistic causal explanation. *Synthese*, 189(3): 415-432.
- Anscombe, G. E. M. (1971). *Causality and determination: an inaugural lecture*. Cambridge University Press.
- Baumgartner, M. and Casini, L. (2017). An abductive theory of constitution. *Philosophy of Science*, 84(2): 214-233.
- Baumgartner, M. and Gebharder, A. (2016). Constitutive relevance, mutual manipulability, and fat-handedness. *The British Journal for the Philosophy of Science*, 67(3).
- Beach, D. and Pedersen, R. B. (2013). *Process Tracing Methods Foundations and Guidelines*. The University of Michigan Press, Ann Arbor.
- Bechtel, W. and Abrahamsen, A. (2005). Explanation: A Mechanist Alternative. *Studies in the History and Philosophy of the Biological and Biomedical Sciences* 36: 421-441.
- Bogen, J. (2005). Regularities and causality; generalizations and causal explanations. *Studies in History and Philosophy of Biological and Biomedical Science*, Vol. 36, No 2.
- Bogen, J. (2008). Causally productive activities. *Studies in the history and philosophy of science*, 39, pages 112-123.
- Cartwright, N. (1989). *Nature's Capacities and Their Measurement*. Oxford University Press, Oxford.
- Cartwright, N. (1995). Ceteris Paribus Laws and Socio-economic Machines. *The Monist*, 78 (3):276-294.
- Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge University Press.
- Cartwright, N. (2002). In Favor of Laws That Are Not Ceteris Paribus after All. *Erkenntnis*, Vol. 57, No. 3: 425-439.
- Cartwright, N. (2009). How to do Things with Causes. *Proceedings and addresses of the American Philosophical Association*, 83 (2), 5-22.
- Couch, M. B. (2011). Mechanisms and constitutive relevance. *Synthese*, 183(3): 375-388.
- Craver, C. (2007). *Explaining the Brain*, Clarendon Press Oxford.
- Craver, C. F. and Bechtel, W. (2007). Top-down causation without top-down causes. *Biology & Philosophy*, 22(4): 547-63.
- Craver, C. F. and Kaiser, M.I. (2013). Mechanisms and Laws: Clarifying the Debate. In H.-K. Chao, S.-T. Chen & R. Millstein (eds.), *Mechanism and Causality in Biology and Economics*. Dordrecht: Springer. pp. 125-145.
- Craver C. and Tabery, J. (2017). Mechanisms in Science. *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2017/entries/science-mechanisms/>.

- Davidson, D. (1963). Actions, Reasons and Causes. *Journal of Philosophy*, 60: 685–700.
- Davidson, D. (1967). Causal Relations. *Journal of Philosophy*, 64: 691–703.
- Flannigan, C., Berry, D., Jarvis, M., and Liddle, R. (2015). *AQA Psychology*, Cheltenham, Glouc: Illuminate Publishing.
- George, A. and Bennett, A. (2004). *Case Studies and Theory Development in the Social Science*, MIT Press.
- Glennan, S. (2005). Rethinking Mechanistic Explanation. *Philosophy of Science* 69, S342-353.
- Glennan, S. (2010). Ephemeral mechanisms and historical explanation. *Erkenntnis*, 72(2): 251-266.
- Glennan, S. (2011). Singular and general causal relations: a mechanist perspective. In Phyllis McKay Illari, Federica Russo & Jon Williamson (eds.), *Causality in the Sciences*. Oxford University Press.
- Glennan, S. (2017). *The New Mechanical Philosophy*. Oxford: Oxford University Press.
- Grunbaum, A. (1963). *Philosophy of science undergraduate lectures*, University of Pittsburgh.
- Harinen, T. (2018). Mutual manipulability and causal inbetweenness. *Synthese*, 195(1): 35-54.
- Hempel, C.G. (1966). *Philosophy of Natural Science*. Englewood-Cliffs, N.J.: Prentice-Hall.
- Hempel, C.G. (1970). *Aspects of Scientific Explanation*. N.Y., N.Y.: The Free Press.
- Hempel, C. G. (2002). Two models of scientific explanation. In Yuri Balashov & Alexander Rosenberg (eds.), *Philosophy of Science: Contemporary Readings*. Routledge. pp. 45—55.
- Illari, P. and Williamson, J. (2012). What is a Mechanism? Thinking about Mechanisms across the Sciences. *European Journal of Philosophy of Science* 2: 119-135.
- Kaiser, M. I. and Krickel, B. (2017). The metaphysics of constitutive mechanistic phenomena. *British Journal for the Philosophy of Science*, 68: 745-47.
- Kaplan, D. M. and Craver, C. F. (2011). The explanatory force of dynamical and mathematical models in neuroscience: a mechanistic perspective. *Philosophy of Science*, Vol. 78, No. 4, pp. 601-627
- Krickel, B. (2017) A regularist approach to mechanistic type-Level explanation. *The British Journal for the Philosophy of Science*, 0(2017), 1-31.
- Krickel, B. (2018). *The mechanical world: the metaphysical commitments of the new mechanistic approach* (Studies in Brain and Mind). Springer.
- Leuridan, B. (2010). Can mechanisms really replace laws of nature? *Philosophy of Science*, Vol. 77, No. 3, pp. 317-340
- Machamer, P., Darden, L. and Craver, C. (2000). Thinking about Mechanisms. *Philosophy of science* 67: 1-25.
- Mill, J. S. (1967). On the Definition of Political Economy. *The Collected Works of John Stuart Mill, Volume IV - Essays on Economics and Society Part I*, ed John Robson, Toronto: University of Toronto Press, 309-40.

Mitchell, S. (1997). Pragmatic laws. *Philosophy of Science, Vol. 64, Supplement. Proceedings of the 1996 Biennial Meetings of the Philosophy of Science Association. Part II: Symposia Papers*, pp. S468-S479.

Pearl, J. (2009). *Causality, Models, Reasoning and Inference* (Second edition). Cambridge University Press.

Pemberton, J. M. (2011). *Integrating Mechanist and Nomological Machine Ontologies to make Sense of What-How-That Evidence*, <https://lse.academia.edu/johnpemberton>.

Pemberton, J.M. & Cartwright, N. (2014). Ceteris paribus laws need machines to generate them. *Erkenntnis Volume 79, Issue 10*, 1745-1758.

Revonsuo, A. (2001). On the Nature of Explanation in the Neurosciences. *Theory and Method in the Neurosciences*, Peter Machamer, Rick Grush and Peter McLaughlin (eds.), University of Pittsburgh Press, 45-69.

Woodward, J. (2014). Scientific explanation. *Stanford Encyclopedia of Philosophy*.
<https://plato.stanford.edu/entries/scientific-explanation/>