

Are higher mechanistic levels causally autonomous?

Peter Fazekas^{1,2,✧} and Gergely Kertész³

¹ *Centre for Philosophical Psychology, University of Antwerp, Belgium*

² *Cognitive Neuroscience Research Unit, CFIN, Aarhus University, Denmark*

³ *Department of Philosophy, Durham University, UK*

✧corresponding author, email: fazekas.peter@gmail.com

Abstract

This paper provides a detailed analysis and explores the prospects of the arguments for higher-level causal autonomy available for the proponents of the mechanistic framework. Three different arguments (a context-based, an organisation-based, and a constraint-based) are distinguished. After clarifying previously raised worries with regard to the first two arguments, the paper focuses on the newest version of the third argument that has recently been revived by William Bechtel. By using Bechtel's own case study, it is shown that not even reference to constraints can establish the causal autonomy of higher mechanistic levels.

1. Introduction

The mechanistic approach aims at accounting for a target phenomenon — typically the behaviour of a higher-level whole — in terms of an underlying mechanism constituted by the organised activities of lower-level parts that jointly produce the very higher-level behaviour in question (Machamer et al. 2000; Bechtel and Abrahamsen 2005; Illari and Williamson 2012). Mechanistic explanations dominate the biological and life sciences (Craver 2007; Craver and Darden 2013), are also important in the physical and engineering sciences (Glennan and Illari, 2018), and are even claimed to be able to contribute to our understanding of causation (Glennan 1996, 2010, 2017). Nevertheless, the ontological commitments and metaphysical implications of the mechanistic framework are far from being clear (Fazekas and Kertész 2011; Soom 2012; Rosenberg 2015; Eronen 2015; Kaiser and Krickel 2016; Krickel 2019).

This paper focuses on the metaphysical implications of the mechanistic approach with regard to the *causal autonomy* of higher levels. Can higher-level wholes be autonomous with respect to the corresponding lower-level mechan-

isms underlying them and producing their characteristic behaviour? According to a wide consensus, they certainly can in an epistemological sense. But what about causal autonomy? Can higher-level wholes possess unique causal powers? This question is hotly debated. On the one hand, it has recently been argued that the causal autonomy of higher-level entities is incompatible with some core commitments of the mechanistic approach (Fazekas and Kertész 2011; Soom 2012; Rosenberg 2015). On the other hand, William Bechtel routinely presents the mechanistic framework as one that is well equipped to ensure the autonomy of higher-level entities even in a strict causal sense (Bechtel 2007, 2008, 2009, 2017a, 2017b; Bechtel and Abrahamsen 2008).

In this paper our aim is to advance this debate by providing a detailed analysis and exploring the prospects of the arguments for higher-level causal autonomy available for the proponents of the mechanistic framework. Mechanists argue for causal autonomy relying on an entangled mixture of different arguments, so the first goal of our endeavour is to disentangle the relevant parts of the literature, and to reconstruct the different arguments in play. We will distinguish three different arguments: a context-based, an organisation-based, and a constraint-based type that are independently motivated by different commitments.

The context-based and organisation-based arguments have played a central role at previous stages of the debate, and have been in the focus of recent objections (Fazekas and Kertész, 2011; Soom, 2012; Rosenberg, 2015). In his recent contributions, Bechtel tries to discredit these objections (Bechtel, 2017a, 2017b), so the second goal of our paper is to clarify how the context-based and organisation-based arguments can be answered.

In his most recent papers Bechtel also emphasises the importance of a third kind, a constraint-based argument for the causal autonomy of higher mechanistic levels (Bechtel, 2017a, 2017b). On the face of it, this line of thought is able to evade existing objections and bestow higher-level entities with unique causal powers. The third and major goal of our paper is to argue that this, however, is not the case. By using Bechtel's own case study, we will show that not even reference to constraints can establish the causal autonomy of higher mechanistic levels.

2. Background: mechanisms, levels and causal autonomy

The mechanistic programme targets a characteristic behaviour of an entity, and aims to describe the mechanism that produces the behaviour in question. Mechanistic explanations proceed via (1) identifying working parts (Bechtel 2008) or

components (Craver 2007) of the target entity, (2) describing the specific spatial and temporal organisation of the parts, and (3) demonstrating that the overall activity of the components organised in such a way is able to produce the behaviour in question (Machamer et al. 2000; Craver 2007; Bechtel 2008).

The spatially and temporally arranged operations of the parts constitute the behaviour of the target entity in the sense that the organised overall activity of the parts exhibit the target behaviour (Craver, 2007). That is, according to the logic of this framework, the mechanistic agenda is to find those components whose joint operations result in the very behaviour that the mechanistic programme aims at accounting for.

The mechanistic approach works with a multi-level picture: the target phenomenon is at a higher level while the entities that together produce the characteristic behaviour of the target are at a lower level (Craver 2007; Bechtel 2008). Levels are defined in terms of the working parts the organised activity of which constitutes the target phenomenon. Entities are at the same level if they are the working parts of the same mechanism: “[i]t is the set of working parts that are organised and whose operations are coordinated to realise the phenomenon of interest that constitute a level” (Bechtel 2008, 146). Similarly, entities are at different levels if one of them (the lower-level entity) is a working part of a mechanism that produces the behaviour characteristic of the other one (the higher-level entity): “X’s ϕ -ing is at a lower mechanistic level than S’s ψ -ing if and only if X’s ϕ -ing is a component in the mechanism for S’s ψ -ing” (Craver 2007, 189).

Since the behaviour of a component can similarly be analysed in terms of the organised activities of still lower level entities, the mechanistic framework presents the world as a nested hierarchy of mechanisms, in which entities residing at a lower level form a mechanism that produces the behaviour of a higher level entity that is a working part of a higher level mechanism, etc.

Those who argue for the causal autonomy of higher levels claim that higher levels are not just causally potent (which would be true even if the causal powers of higher-level entities were also the powers of certain lower-level entities; see e.g. the subset view: Shoemaker, 2007) but also possess unique causal powers. As it has lately been emphasised, such an argument for causal autonomy needs to start with establishing the claim that lower levels are causally not closed (Hendry 2010). William Bechtel (2017a, 2017b) has recently taken up the challenge and tried to demonstrate that the way higher level constraints and boundary conditions work disproves the causal closedness of lower levels, renders the mechanism as a whole not just causally potent, but also caus-

ally autonomous, and helps evade the causal exclusion argument (Kim 1998). Note that to support such a causal autonomy claim one needs to demonstrate that the causal influence a higher-level entity exerts to determine the unfolding of some lower-level events uniquely belongs to the higher-level, ie. it is not exerted by any lower-level entity.

3. Arguments for higher-level causal autonomy

Over the last two decades many attempts to argue for the causal autonomy of higher levels have been published, often presenting an entangled web of three different arguments.

3.1 The argument from context

An important consequence of how the mechanistic framework defines the criteria of being at the same level and being at different levels (see Sec. 2) is that levels of mechanisms are *local*: they are well-defined only within a given compositional hierarchy. If two entities are not working parts of the same mechanism, then there is no meaningful way to address the question whether they are at the same level (Craver 2007, 192; Bechtel 2008, 147). Due to this locality, lower levels are never causally closed, since lower levels are never extended beyond the set of entities that together constitute a mechanism, and thus they are always restricted and partial. So within the mechanistic framework there are no comprehensive lower levels that are causally complete and closed (Bechtel 2008, 148). Local levels restricted to the constituents of a mechanism typically lack resources to account for effects exerted from outside the mechanism, i.e. from the *context* in which the higher-level whole is embedded (Bechtel 2007, 183; 2008, 152). As a consequence, such lower level effects will have only higher level causes.

3.2 The argument from organisation

The spatial arrangement of the entities forming a mechanism and the temporal organisation of their activities are crucial determinants of a mechanism. However, just as resources necessary to account for contextual effects, resources required to account for arrangement and organisation are also unavailable at the lower level of the parts. The spatial and temporal structure of the parts is independent of their behaviour: parts conforming to the same laws and producing identical behavioral repertoire in isolation can nevertheless be organised into very different structures (Bechtel 2007, 183). This organisation is imposed upon

the parts by the higher level whole (Bechtel 2008, 150; 2009, 554-57), and thus is a manifestation of the unique causal powers of higher levels.

3.3 The argument from constraints

The spatial and temporal organisation of the parts constrains the behaviour of the parts: modes of organisation imposed on the constituents restrict how they can interact with each other (Bechtel 2009, 555-57). Similarly, the functioning of the mechanism as a whole in its higher-level context impinges specific conditions upon the mechanism that, at least partly, dictate how the parts can operate (Bechtel 2008, 240; 2009, 557-59). So by imposing a structure on the parts and by interacting with its environment the higher level whole constrains the behaviour of the parts — which is an extra causal influence on what happens inside the mechanism exerted by higher level entities, and thus a further source of the causal autonomy of higher levels (Bechtel 2017a, 271).

4. Clarifying the answers to the arguments from context and organisation

Recent criticisms of the idea that higher mechanistic levels are causally autonomous provided reasons to resist the arguments from context and organisation. It has been argued that the way the mechanistic framework thinks about constitution, causation and levels are incompatible with the main claims of these arguments.

4.1 Identifying higher and lower level causal roles

It is a fundamental tenet of the mechanistic framework that the organised activity of the constituent parts of a mechanism produces the very behaviour that characterises the target phenomenon, since this is the requirement that ensures the success of explaining a higher-level phenomenon mechanistically in terms of the behaviour and organisation of certain lower-level entities. If the behaviour produced by the organised activity of the lower-level entities was not identical to the characteristic behaviour of the target phenomenon, then the mechanistic story would clearly not be an account of the target phenomenon (Fazekas and Kertész 2011; Soom 2012; Rosenberg 2015). That is, the mechanistic framework is inherently committed to *identifying the causal roles* a whole plays at the higher-level with the causal roles the organised activity of the parts plays at the lower-level. So, even if causal processes at different levels look different, they are not different—in fact, the corresponding ones must be identical (Bechtel 2008, 2009; Fazekas and Kertész 2011).

Note that this picture is not eliminativist, for higher-level entities are token-identical with sets of lower-level ones. It is not epiphenomenalist either, for higher-level entities do possess causal powers — the same causal powers that are also possessed by certain spatially and temporally structured sets of lower level entities (see Fazekas and Kertész 2011; Soom 2012). This latter point needs special emphasis, given that Bechtel interprets his opponents as if they were arguing for the unjustified extreme view, which is “highly reductionist”, “represents all activity at one lowest level” (Bechtel 2017a, 269) and renders “higher levels epiphenomenal” (Bechtel 2017a, 262; for such a view see Rosenberg 2015). Contrary to this, the picture advocated here is compatible with the usefulness of higher-level enquiries (as the only sources of higher level descriptions), and the importance and significance of higher-level accounts (as the only sources of certain generalisations, as say, in the case of multiple realisability); i.e. it is compatible with the epistemic or explanatory autonomy of higher levels. What it is incompatible with is the view Bechtel wants to argue for: the *causal* autonomy of higher levels.

4.2 Answering the argument from context

Another commitment explicitly endorsed by proponents of the mechanistic framework is that causation is an intra-level phenomenon: causal links do not span between different levels, they are confined to single levels. As Craver and Bechtel put it: “[t]here are no causal interactions beyond those at a level” (2007, 561; see also Craver 2007; Bechtel 2008, 2017; Fazekas and Kertész 2011). Note, however, that this commitment to causation as an exclusively intra-level phenomenon is in tension with the claim that levels are local. Due to locality, even if it is possible to analyse two entities that are interacting, causally connected parts of a mechanism in terms of lower-level sub-mechanisms, it is not possible to decide whether the lower-level parts constituting the sub-mechanism responsible for the behaviour of one of the entities are at the same level as the lower-level parts constituting the other sub-mechanism (Craver 2007).

The problem with this consequence of locality is the following: (i) in accordance with the logic of mechanistic explanations, the organised activity of the parts constituting the sub-mechanisms produce the very behaviours that are characteristic of the entities analysed; (ii) if an interaction between the two entities is part of this behaviour then there will be an interaction between the two sub-mechanisms as well; (iii) and if causality is strictly an intra-level phenomenon then the causal connection between the two sub-mechanisms entails that the

two sub-mechanisms are at the same level (Fazekas and Kertész 2011; Eronen 2015).

That is *lower levels can be extended* beyond the original scope of mechanistic enquiry. Consequently, lower levels are not necessarily restricted and incomplete, and thus there is no principled reason to think that they are causally not closed. In particular, note that the higher-level context that plays a crucial role in the ‘argument from context’ is the higher-level environment of the entities targeted by mechanistic decomposition. This environment consists of those higher-level entities that interact with the entity under scrutiny. Since lower levels can be extended exactly along the lower-level counterparts of such higher-level connections, contrary to the claim of the ‘argument from context’, resources are very much available at the lower-level to account for the effects exerted by the context of the mechanism (Fazekas and Kertész 2011, Soom 2012).

4.3 Answering the argument from organisation

Bechtel explicitly argues that how the components are spatially, temporally, and relationally organised go beyond the account of the parts and their operations. However, a lower-level account is not restricted to the (intrinsic) characteristics of the parts and the behavioural repertoire they produce in isolation. Thinking so would simply misidentify the supervenience base of the higher level entity under scrutiny. The organisation of the lower-level parts crucially determines what kind of behaviour the overall mechanism produces. So a lower-level account must include information about the spatial and temporal structure of the parts, and their interactions. In fact, the organisation of parts can only be accounted for at the lower-level, since it is the lower-level methodology that can target lower-level entities and thus can uncover their organisation, and it is the lower-level terminology that is apt for describing the parts and their structure. Note that these are criteria that Bechtel himself proposes as ways to characterise higher and lower levels: he claims that at higher and lower levels different experimental and explanatory strategies and different vocabularies are needed (Bechtel 2007, 185; 2008, 155-57). That is, since spatial, temporal and relational facts of lower-level entities are discovered by lower-level enquiries, and are characterised by lower-level vocabularies, by Bechtel’s own standards they are very much integral parts of lower-level accounts (Fazekas and Kertész 2011).

5. Answering the argument from constraints

Bechtel (2017a) acknowledges that there is room for debate with regard to his approach to the autonomy of higher levels. As he admits, he and Craver have

“failed to bring out in what sense higher levels are involved in producing [...] effects at the lower-level” (Bechtel 2017a, 254). Bechtel aims to clarify this issue by emphasising the role constraints play in shaping lower-level causal processes, i.e. by providing a refined version of the ‘argument from constraints’.

Bechtel claims that parts embedded in a mechanism behave differently than in isolation because of the restrictions imposed upon them (Bechtel 2017a, 271). These restrictions limit the behavioural repertoire of the parts by radically reducing the degrees of freedom available to them, thereby forcing them into a narrow behavioural space. Bechtel emphasises those cases in which restrictions are generated by feedback loops. Feedback makes the behaviour of the individual components sensitive to the actual state of the whole mechanism that they are embedded in. Since such a state can change in accordance with a characteristic dynamics, the observable behaviour of the parts can also change with time resulting in surprising patterns of different effects evoked by the same input (Bechtel 2017a, 272-273).

Note that on the face of it this strategy has the potential to avoid the worries raised with regard to the arguments from context and organisation. Even if the causal roles of higher-level wholes (the behaviours produced as responses to specific influences) are identical to the causal roles of sets of interacting lower-level parts constituting a mechanism, and even if lower levels are extendable along the causal connections of the lower-level entities constituting the mechanism to non-constituents outside the mechanism boundary, constraints can still restrict what can happen *inside* the mechanism. If such constraining effects can be attributed to no lower-level entities but only to the higher-level whole, then the ‘argument from constraints’ goes through, and the higher-level proves to be causally autonomous. (Here we follow Bechtel who argues that constraints are causal. For a different view, see, for example: Paoletti and Orilia 2017, 4-7.)

5.1 A case study: the circadian clock mechanism

Bechtel’s own case-study to support the ‘argument from constraints’ is about the intra-cellular molecular ‘clock mechanism’ that is responsible for circadian rhythmicity — cyclic oscillations of daily behaviours and physiological functions. The molecular clock mechanism partly consists of a transcriptional-translational feedback loop involving so-called ‘clock genes’ *Per* and *Cry* and their protein products PER and CRY. Inside the nucleus, transcription factors BMAL1 and CLOCK drive the transcription of the genes *Per* and *Cry* to RNAs, which then are decoded by ribosomes to produce proteins PER and CRY. After a translocation back to the nucleus, PER and CRY inhibit the transcriptional

effects of BMAL1 and CLOCK through direct protein-protein interactions. Via this negative feedback loop PER and CRY autoregulate their own transcription, which results in a periodic increase and decrease in their concentration, and defines the phases of the oscillation of the clock mechanism. During the night PER and CRY concentration is high, while it is low during the day such that it peaks in the middle of the night after which it decreases and bottoms out in the middle of the day. Correspondingly, the level of *Per* and *Cry* transcription is low during the night, increases at dawn as PER and CRY concentration falls, at its maximum during the day and then decreases at dusk (the presentation here is simplified; for full details see Rosenwasser and Turek 2017). That is, the transcription of *Per* and *Cry*, which is a component of the mechanism, depends on the phase of the oscillation, and thus is sensitive to the actual state of the mechanism as a whole (Bechtel 2017a, 257). The constraining effects of the higher level whole on the behaviour of the components can most clearly be seen, Bechtel argues, in the different effects of light exposure on *Per* transcription at different times of the day: while light input has no effect during daytime, at dusk it delays, whereas at dawn it advances the phase of the oscillation (Bechtel 2017a, 267).

Bechtel's claim here is that characteristics of the lower-level mechanism, like the behaviour of certain parts (the level of transcription of *Per* and *Cry*), and how the organised activity of the parts processes a given input (exposure to light) are determined by the actual state of the higher-level whole (the phase of the oscillation). Recall that Bechtel uses this example to support his 'argument from constraints' against the causal closure of the lower-level and for the causal autonomy of the higher-level (see e.g. Bechtel 2017a, 272). So the claim that he really needs is that the way in which the higher-level whole determines the unfolding of lower-level events is via a causal influence, and that no lower-level entity exerts the same causal influence. Does Bechtel's own example support this claim?

5.2 Constraining effects are exerted by lower-level entities

Bechtel crucially relies on the phase of the oscillation of the circadian clock mechanism as the determinant of how lower-level processes unfold. The phase of the oscillation, however, is not a characteristic of some higher-level whole that has no counterpart at the lower level. On the contrary, the molecular description that characterises the clock mechanism can and does define the phases in purely lower-level (molecular) terms: as the periodic changes in the concentrations of PER and CRY (see above, also Rosenwasser and Turek 2017, 352-55).

So the phase of the oscillation is not some higher-level factor that causes changes in the level of transcription of *Per* and *Cry* in a way such that the relevant causal power is not possessed by any lower-level entity. It is the changes in the concentrations of PER and CRY that cause changes in the level of transcription of *Per* and *Cry* — via the direct protein-protein interactions with BMAL1 and CLOCK inhibiting their drive of the transcription of *Per* and *Cry*. Similarly, it is not the phase of the oscillation, as a higher-level entity per se, that constrains how the parts of the mechanism change their behaviour as a response to a light signal, but the actual level of concentrations of PER and CRY (see Bechtel 2017a, 267-68).

Interestingly, Bechtel is in agreement with us with regard to the presence and importance of these lower-level interactions. He acknowledges that “the phase of the oscillator at a time *just is* the concentrations of PER, CRY” (Bechtel 2017a, 257, our emphasis). And he also acknowledges that “[a]s a result of the interconnectivity of the parts, especially the feedback loops, the [...] mechanism functions as a unit, with the operations of the individual parts of the mechanism *determined by other parts of the mechanism*” (Bechtel 2017a, 268, our emphasis). Nevertheless, he thinks that these interactions are compatible with the causal autonomy of higher levels (Bechtel 2017a, 272). But they are not. These interactions, by ensuring that the lower-level effects in question have lower-level causes, causally close the lower-level with regard to the behaviours under scrutiny, and leave no room for the causal autonomy of the higher-level.

Per, *Cry*, PER, CRY, BMAL1, CLOCK, etc. are interacting working parts of the clock mechanism, and as such — by the standards of Bechtel’s own definition that puts working parts of a mechanism at a lower level than the whole the behaviour of which is produced by the mechanism (Bechtel 2008; see also Craver 2007) — they are lower-level entities. Furthermore, they, their interactions, their spatial arrangement and the temporal organisation of their activities are all studied by lower-level methodologies, and characterised by lower-level vocabularies (see Sec. 4.3). Is there contextual information that cannot be captured at this level? No: the relevant contextual information concerns the effect of light on the phase of the clock mechanisms — but this is also defined in terms of the lower level, as the boosting effect of light on *Per* and *Cry* transcription. What is left? Nothing. All information is there at the lower level, and all causal influences can be exerted by lower-level entities. The higher level possesses no unique causal powers that couldn’t be attributed to some lower-level factor. Therefore, the higher-level is not causally autonomous.

Bechtel is right that feedback loops are important. But they are important as parts of lower-level interactions. Having an account of how feedback loops actually work reveals how lower-level processes impose complex constraints upon their own unfolding, and thus, instead of disproving, it contributes to appreciating the causal closedness of the lower level.

6. Conclusion: higher mechanistic levels are not causally autonomous

Here we distinguished three different arguments for the causal autonomy of higher mechanistic levels: the arguments from context, organisation and constraints. Upon closer reflection, it is evident that the constraining effects that restrict and determine the behaviours of lower-level parts are brought about by the interactions between lower-level entities. Parts embedded in a mechanism behave differently than in isolation because they are in constant interactions with other parts of the mechanism and also with further entities external to the original mechanism. These interactions internal and external to a given mechanism are what define the organisation of the parts and the context of the mechanism, respectively, thereby encoding both organisation and context based constraints at the lower-level. The effects of context, organisation and constraints can all be accounted for in terms of the causal influences of lower level entities and activities. That is, within the mechanistic framework, the causal autonomy of higher levels cannot be established.

References

- Bechtel, William. 2007. "Reducing psychology while maintaining its autonomy via mechanistic explanation." In *The Matter of the Mind: Philosophical Essays on Psychology, Neuroscience and Reduction*, ed. Maurice Schouten and Huib Looren de Jong, 172-198. Oxford: Basil Blackwell.
- . 2008. *Mental mechanisms: philosophical perspectives on cognitive neuroscience*. New York: Psychology Press.
- . 2009. "Looking down, around, and up: Mechanistic explanation in psychology." *Philosophical Psychology*, 22:543-564.
- . 2017a: "Explicating top-down causation using networks and dynamics." *Philosophy of Science*, 84:253-274.
- . 2017b: "Top-down causation in biology and neuroscience: Control hierarchies." In *Philosophical and scientific perspectives on downward causation*, ed. Michele Paolini Paoletti and Francesco Orilia. London: Routledge.
- Bechtel, William, and Adele Abrahamsen. 2008. "From reduction back to higher levels." In *Proceedings of the 30th Annual Conference of the Cognitive Science Soci-*

- ety, ed. B. C. Love, K. McRae and V. M. Sloutsky, 559-564. Cognitive Science Society.
- Craver, Carl. 2007. *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. New York: Oxford University Press.
- Craver, Carl, and William Bechtel. 2007. "Top-down causation without top-down causes." *Biology and Philosophy*, 22:547-563.
- Craver, Carl, and Lindley Darden. 2013: *In Search of Mechanisms. Discoveries across the Life Sciences*. Chicago: University of Chicago Press.
- Eronen, Markus. 2015. "Levels of organization: a deflationary account." *Biology and Philosophy*, 30:39-58.
- Fazekas, Peter, and Gergely Kertész. 2011: "Causation at different levels: tracking the commitments of mechanistic explanation." *Biology & Philosophy*, 26:365-383.
- Glennan, Stuart. 1996. "Mechanisms and the nature of causation." *Erkenntnis*, 44:49-71.
- . 2010. "Mechanisms, causes, and the layered model of the world." *Philosophy and Phenomenological Research*, 81:362-381.
- . 2017. *The new mechanical philosophy*. Oxford: Oxford University Press.
- Glennan, Stuart, and Phyllis Illari. 2018. *The Routledge Handbook of Mechanisms and Mechanical Philosophy*. New York: Routledge.
- Illari, Phyllis, and Jon Williamson. 2012. "What is a mechanism? Thinking about mechanisms across the sciences." *European Journal for Philosophy of Science*, 2:119-135.
- Kaiser, Marie I., and Beate Krickel. 2016: "The metaphysics of constitutive mechanistic phenomena." *The British Journal for the Philosophy of Science*, online first, 1-35.
- Kim, Jaegwon. 1998. *Mind in a physical world*. Cambridge, MA: MIT Press.
- Krickel, Beate. 2019. *The Mechanical World – The Metaphysical Commitments of the New Mechanistic Approach*. Springer.
- Machamer, Peter, Lindley Darden, and Carl Craver. 2000. "Thinking about Mechanisms." *Philosophy of Science*, 67:1-25.
- Paoletti, Michele Paolini and Francesco Orilia. 2017. *Philosophical and Scientific Perspectives on Downward Causation*. New York, NY: Routledge.
- Rosenberg, Alex. 2015. "Making mechanism interesting." *Synthese*, Electronically published May 30. doi:10.1007/s11229-015-0713-5.
- Rosenwasser, Alan M. and Fred W. Turek. 2017. "Physiology of the Mammalian Circadian System". In *Principles and Practice of Sleep Medicine*, ed. Meir. H. Kryger, Thomas Roth, and William C. Dement, 351-361. Elsevier Health Sciences.
- Shoemaker, Sydney. 2007. *Physical Realization*. Oxford: Oxford University Press.
- Soom, Patrice. 2012. "Mechanisms, determination and the metaphysics of neuroscience." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 43:655-664.