

Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

Studies in History and Philosophy of Biological and Biomedical Sciences

journal homepage: www.elsevier.com/locate/shpsc

A weakened mechanism is still a mechanism: On the causal role of absences in mechanistic explanation



Alexander Mebius

Royal Institute of Technology (KTH), Department of Philosophy and History of Technology, Teknikringen 78 B, 100 44 Stockholm, Sweden

ARTICLE INFO

Article history:

Received 20 August 2013
 Received in revised form 5 November 2013
 Available online 28 November 2013

Keywords:

Causation by absence
 Mechanisms
 Causal production
 Attenuation
 Mechanism failure
 Benjamin Barros

ABSTRACT

Much contemporary debate on the nature of mechanisms centers on the issue of modulating negative causes. One type of negative causability, which I refer to as “causation by absence,” appears difficult to incorporate into modern accounts of mechanistic explanation. This paper argues that a recent attempt to resolve this problem, proposed by Benjamin Barros, requires improvement as it overlooks the fact that not all absences qualify as sources of mechanism *failure*. I suggest that there are a number of additional types of effects caused by absences that need to be incorporated to account for the diversity of causal connections in the biological sciences. Furthermore, it is argued that recognizing natural variability in mechanisms, such as attenuation, leads to some interesting line-drawing issues for contemporary philosophy of mechanisms.

© 2013 The Author. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/3.0/>).

When citing this paper, please use the full journal title *Studies in History and Philosophy of Biological and Biomedical Sciences*

1. Introduction

The specification of negative causes is an integral part of molecular biology and neurobiology. Terms such as antagonists, blockers, repressors, depressors, gates, and gene-knockouts, for example, exemplify instances of negative causation. It is widely agreed that these terms play a fundamental role in mechanistic explanation (Craver, 2007; Machamer, 2004). There are several types of causal claims that might be regarded as “negative”, including preventions/interferences, omissions/absences, disconnections, and mixtures of these. In this paper, I scrutinize the putative problem of causation by absence as it occurs in the context of mechanistic explanation.

According to one of the most prominent accounts of mechanistic explanation, advocated by the team of Peter Machamer, Lindley Darden and Carl Craver (2000), mechanisms are specified through descriptions of their entities and activities that are organized so as to causally produce some phenomenon. Since the mechanistic production view of causation requires some sort of transfer of energy

or force from one entity to another, an absence cannot participate in such a transfer. They lack the appropriate “oomph”, which is the physical connection between causes and their effects (Beebe, 2004; Dowe, 2004; Schaffer, 2004, 2005). Therefore, it is hard to see how there can be a physically dependent relationship (i.e. oomph between C and E) in terms of production between a nothing and a something.¹ Barros (2013) recently suggested a solution to this problem that relies on interpreting absences as causes of mechanism failure.

This paper argues that resolving this problem by categorizing absences as instances of mechanistic failure is unsatisfactory. It is argued that, at least when it comes to the biological sciences, one can speak of *types* of functionality, and *degrees* thereof, pertaining to the same biological mechanism affected by an absence. As such, these instances do not correspond well with the notion of mechanism failure.

The next section of this paper discusses the notion of causal productivity as it is commonly conceived of in leading accounts

E-mail address: alenik@kth.se

¹ Note that the problem of absence is not tied to mechanistic explanation *per se* but is a problem about causation and causal explanation in general. Nevertheless, the discussion here is confined to the problem as it appears in the mechanistic context.

of mechanistic explanation. I then provide some examples of absences and suggest why they are essential for explaining biological mechanisms in sciences such as neurobiology and physiology. Section 4 considers the recent attempt by Benjamin Barros (2013) to characterize absences as instances of mechanism failure. Here I argue that Barros' solution is not complete since it leaves out several important types of causes and effects associated with absences in the biological sciences, notably the effects of attenuation. Section 5 presents a more nuanced approach for incorporating absences. The penultimate section presents the consequences of mechanism attenuation for the philosophy of mechanisms. The final section concludes the paper.

2. Causally productive mechanisms

This section considers the role of causal production in describing mechanisms. More specifically, I suggest why incorporating absences with mechanistic concepts of productivity have proved elusive.

Philosophers of science have, in more recent years, offered a plethora of mechanistic characterizations of causality and causal explanation (e.g. Bechtel & Abrahamsen, 2005; Craver, 2007; Glennan, 1996, 2002; Illari & Williamson, 2012; Machamer, Darden, & Craver, 2000; Torres, 2009; Woodward, 2011). Although there are considerable dissimilarities between these approaches, they are extraneous to the points made in this paper. So, for clarity of exposition, the main focus here is on the Machamer et al. (2000) account. These authors characterize mechanisms as follows:

Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions. (Machamer et al., 2000, p. 3)

At the heart of Machamer, Darden and Craver's analysis is the idea of physical dependency between the organizational features responsible for producing, maintaining and underlying phenomena. Notably, these features can only be accounted for by ontic referents. So, for example, in explaining how a neuron's membrane potential depolarizes during an action potential, the ontic condition confines the citing of *explanans* to physical objects, such as the increase in permeability to Na^+ (Craver, 2007). This means that only entities and activities physically present in mechanisms, such as Na^+ , qualify as component parts of mechanisms and explainers thereof.

For this reason, modern mechanists place significance in the ontic structure that mechanisms describe by seeking to characterize the link between the causal producers of change. This strategy turns on connecting mechanistic models to physical structures, i.e. mechanisms responsible for producing the phenomenon of interest. The more specific notion of ontic adequacy requires that the elements described in the model map the elements found in the mechanism. Only in this respect, by describing the links involved in the operation of mechanisms, do mechanisms become informative in the ontic sense (Craver, 2006; Kaplan, 2011).²

However, Machamer, Darden and Craver's way of framing causal interaction as physically continuous processes becomes problematic when implementing absences as causes and effects in descriptions of mechanisms. It would appear that absences resist any stronger ontological reading because, as Jonathan Schaffer puts it, "absence causation is *metaphysically abhorrent*. When the gardener does not water my flowers, there is no energy-momentum flow or other physical process connecting this absence (wherever located, if at all) to the wilting flowers. Absences impart no 'oomph'." (Schaffer, 2005, p. 300, emphasis in original). Hence, it

appears that the problem with fitting absences into mechanistic explanations of productivity is that absences violate the (mechanistic) condition of physical connectivity since they impart no "oomph".

It is my hope that this section has explained why the concept of causal productivity, as it is currently conceived of in terms of physical continuous processes, appears to be incompatible with the idea of absences as causes and causal explainers of mechanisms. In the discussion that follows, I argue that absences are an essential part of mechanistic explanation by referencing examples of their role as causes and effects in the biomedical sciences and beyond. The idea is that the problem of incorporating absences has to be resolved to obtain a more complete understanding of the way mechanisms and mechanistic explanations are essential concepts in these sciences.

3. Causation by absence in the biological sciences

In this section I explore absences in scientific practices where they are cited most frequently. I also illustrate two features of absences that deserve emphasis. First, that cases of causation by absence, such as in Helen Beebee's example, "I killed the plant by not watering it" (Beebee, 2004), display no physical continuity between Beebee's omission to water the plant and its subsequent death. Second, I suggest that there is a slight yet important difference between describing the *lack of* calcium as causing stunted plant growth, and describing the (total) absence of calcium as causing the death of the plant's terminal buds.

It is commonly conceived that absence causation is the notion that treats the non-existence of an event as a cause/effect. Of course, there are innumerable examples to choose from for the purpose of illustrating causes and effects by absences, a fact that has been duly noted by many philosophers, especially Schaffer: "What causes scurvy is an absence of vitamin C, what causes rickets is an absence of vitamin D, what causes diabetes mellitus is an absence of insulin..." (Schaffer, 2004, p. 202). I center the discussion here on two interesting examples from physiology and neurophysiology.

The first example considers lactose intolerance, which is a disorder affecting the digestion of lactose. Lactose is a disaccharide composed of glucose and galactose and is replete in milk and milk-based food products. The inability to absorb lactose is most probably caused by a deficiency of the enzyme lactase (Vesa, Marteau, & Korpela, 2000). This intestinal enzyme is responsible for hydrolyzing lactose into its constituent monosaccharides, glucose and galactose. Monosaccharides such as glucose and galactose are simple carbohydrates that can be absorbed in the digestive tract. Lactase makes digestion possible by breaking down more complex carbohydrates, such as lactose, because disaccharides like lactose and sucrose are too large to be absorbed (i.e. transported from the intestinal lumen across the epithelium into the bloodstream).

Lactose intolerance is one of several digestive disorders that might be caused by a deficiency in the lactase gene or induced by environmental factors. Such deficiencies are caused by an absence. In this case, the absence can be understood as either a piecemeal lack of something, such as lactase or calcium, for example, or it can be understood as the complete non-existence of something, such as a regulatory gene. Seemingly, biology and biomedicine has an abundance of each type. It goes without saying that biomedical scientists investigating the causes of these diseases would be vexed by the view that absences are not genuine causes, since reasoning about absences underpins most of their work.

² To be clear, I am not conflating the ontic view with the physical process view here. One could be committed to an ontic explanation without being committed to a physical process account of causation.

In any event, it is easy to notice the transitivity common in many mechanistic processes. For instance, the very common chain of events when the absent lactase gene (absent from birth) leads to lactose intolerance, which in turn leads to decreased milk (calcium) consumption, which may then lead to nutrition-related diseases such as osteomalacia and rickets, as regularly effected by calcium deficiency. This example illustrates that a complicated series of interconnected events involving effects/causes due to absences are common in contemporary physiology.

I turn now to the second example. In neurobiology, the last couple of decades have witnessed some very exciting research on the biochemistry underlying the formation of long-term memory.³ The phenomenon of long-term potentiation is considered a useful model for studying synaptic changes that underlie memory and learning. It is commonly understood that the *N*-methyl-D-aspartate (NMDA) receptor⁴ has a pivotal role in the mechanism of long-term potentiation.

NMDA is a structural analog of glutamate, which is a negatively charged amino acid and a constituent of protein. In the central nervous system, glutamate (i.e. glutamic acid) acts as an excitatory neurotransmitter that binds to glutamate receptors, such as the NMDA receptor, and forms a channel through the membrane of postsynaptic neurons. As NMDA binds to the NMDA receptor, the ion channel opens by discharging the Mg^{2+} ion blocking the channel. This allows the intracellular build-up of calcium (Ca^{2+}) ions to diffuse through the neuronal membrane via the ion-channels into the cell. Since the Mg^{2+} ions prevent Ca^{2+} ions from entering the ion-channels, the absence of Mg^{2+} is an intermediate cause of the influx of Ca^{2+} in the process, ultimately leading to the strengthening of the synapse (long-term potentiation).

This example of long-term potentiation, which has become the pet example of philosophers of science, illustrates how an absence can effectuate the productivity of a causal process without there being a physical connection between every cause and effect in the mechanistic chain of events along the way. Evidently, there is no physical interaction taking place between the Mg^{2+} and Ca^{2+} ions in the mechanism of long-term potentiation.

The lesson from these examples is to recognize how biological and neurobiological practice treats absences as causal producers of change. In this respect, the integration of absences becomes crucial for making sense of their roles as causes and causal explainers of mechanisms in scientific practice.

4. Absences as instances of mechanism failure

One proposed way of resolving the incorporation problem is by treating absences as causes of mechanism failure. The literature on causation provides many accounts of causal failures, usually represented as deviations from default values (Menziez, 2009). By incorporating these ideas in the context of mechanistic explanation, Benjamin Barros (2013) proposes an interesting approach that considers a mechanism's failure to operate, as caused by the absence of background conditions and/or entities required for that mechanism:

Under this approach, an operational mechanism provides a baseline with which to contrast the mechanism's failure, and the absence provides the causal explanation for differing outcomes. Mechanism failure may be caused by the absence of a background condition required for the mechanism's operation,

or it may be caused by the absence of a part or entity from the mechanism. (Barros, 2013, p. 466)

Barros offers two ways to describe absences as causes of mechanism failure.⁵ The first is by citing one or several missing background conditions as *explanans*. Barros illustrates this idea with an example involving a match-striking mechanism designed to operate in the normal earth atmosphere.

It follows that as long as the device operates in the normal earth atmosphere it will, *ceteris paribus*, maintain its operational ability: that of igniting a match. However, if the match-striking device is instead placed in a box containing nothing but pure nitrogen, the device would fail to operate (i.e. the match would not ignite). The reason why the match-striking device fails to ignite the match in the latter case is due to the absence of oxygen. This information is obtained by comparing the first instance, where the match-striking device operates normally, with the latter instance, where the match-striking device fails to operate. Hence, the operational failure indicates that oxygen is a background condition required for the match to ignite. The absence of oxygen can therefore be cited as explaining the match-striking device's failure to operate.

The second way of explaining a mechanism's failure to operate is by citing a part or entity that is missing from that mechanism. Barros uses Schaffer's well-known gun firing example to illustrate his point (see Schaffer, 2005). In short, if a certain part is missing from the gun firing mechanism, the mechanism might fail to operate by reference to the absence of that part, e.g. a spring or a level. The state of the mechanism's failure is explained by comparing it with the state of a functional mechanism, that is, by contrasting the non-operational structure with the operational structure.

It does not seem unreasonable to assert, as Barros does, that the difference in outcome can at times be explained by the absence of necessary conditions as required for the mechanism to operate. However, it might be a bit too presumptuous to conclude that, "[a]bsences [...] can be best understood as a part of a causal explanation of mechanism failure" (Barros, 2013, p. 468). Let me explain why.

The notion of "mechanism failure" implies a complete loss of a mechanism's operational capacity (i.e. its ability to function). However, loss of function in a mechanism is essentially *gradual*. Consider, for example, instances of functional decay such as Parkinson's disease, where the absence (deficiency) of dopamine-generating cells in the substantia nigra leads to a gradual loss (atrophy) of memory and other cognitive abilities. The process of memory decay is described as "gradual" because the production of dopamine is severed and diminishing, but it is not entirely disrupted. Although an absence is responsible for the decaying effect on memory, the effect in question does not satisfy Barros' criteria for failure since there is no disruption in operation.

I suggest that the phenomena under study, which could be called "the effects on mechanism dynamics caused by absences," is transitory and varied. As such, the absence of an entity or activity can both strengthen and weaken a mechanism. In addition, we saw that the failure of dopamine-generating cells to produce dopamine leads to the absence of dopamine, which, in turn, affects cognition.

5. Absences as instances of mechanism attenuation and mechanism enhancement

Under Barros' conditions, absences are useful for distinguishing between operational and non-operational mechanisms. In more

³ Numerous experiments with transgenic mice as models of spatial learning have gained the attention of philosophers in debates about reduction and reductive explanation. The methodology inherent in this research turns on creating genetically engineered mutant mice lacking the expression of certain genes and then tracing the effects of the absence on cognition, see, e.g. John Bickle's (2006) paper.

⁴ NMDA receptors are ligand-gated ion channels acting as mechanisms for the modulation of synaptic plasticity and memory.

⁵ Ned Hall makes a similar point regarding the notion of failure by suggesting that an absence can be understood as describing "the failure of *c* to occur" (Hall, 2004, p. 248).

indefinite situations, however, the gain achieved by conceptualizing absences as leading to mechanism failure is not as clear-cut. It is my view that the notion of causation by absence has the advantage of going beyond merely characterizing and explaining all-or-nothing phenomena. If we were interested, for example, in whether absences add to or diminish certain operational features of any given mechanism, we would have to identify additional causal arrangements that are caused by absences. In doing so, we need to expand on the idea that absences are the causal explainers of mechanism failure by recognizing additional modes of effect caused by absences.

In neurobiology, as in the biological and biomedical sciences in general, absences can lead to *attenuation*, a phenomenon that is very different from failure or malfunction. As mentioned previously, mechanisms negatively affected by an absence usually entail the debilitation of a mechanism, that is, the depletion of strength or exhaustion of a mechanism, but rarely a complete operational failure. Such degenerating effects are evident in cases of clinical depression, where the normal function of synaptic transmission is affected by the absence of dopamine, and hence the mechanism's operation is weakened, but nonetheless operational. Research reports in bioscience and bioengineering journals provide many cases of attenuation, evidenced by article titles such as “Acetaminophen attenuates dopamine neuron degeneration in animal models of Parkinson's disease” (Locke, Fox, Caldwell, & Caldwell, 2008) and “Sensorimotor attenuation by central motor command signals in the absence of movement” (Voss, Ingram, Haggard, & Wolpert, 2006). I would even go so far as to suggest that instances of mechanism attenuation are more common than instances of mechanism failure, at least in the biological sciences. For instance, Type 2 diabetes, the form that occurs when the body produces insufficient amounts of insulin, is considerably more common than Type 1 diabetes, the form that occurs when the body does not produce any insulin at all (Evans, Thornton, Chalmers, & Glasziou, 2011). The effects of the absences in question are therefore best understood as part of the causal explanation for mechanism attenuation.

To explain the idea of attenuation, consider what electrical engineers call an attenuator device. This electronic device reduces the power of a signal *without* distorting its (fundamental) waveform. Similarly, an amplifier is a device that works in the opposite manner; increasing the power of a signal. These two devices provide either loss or gain in signal amplitude. Similar to the range of possible variations in signal amplitudes, the possible effects of absences come in a continuous scale of varying degrees of effects, ranging from a small lack of, to a complete (total) absence of, entities or activities underlying mechanistic phenomena. In contrast to malfunction, examples of attenuation dissipate power, without distorting the (fundamental) signal.

Recall the gun firing example above in which Barros suggests that the absence of a part or entity from the operational baseline of the mechanism leads to mechanism failure. The following example aims to show that even the absence of a baseline component of a mechanism does not necessarily entail failure, or even attenuation.

Neuropharmacological agents, known as serotonin reuptake inhibitors, decrease the breakdown of the neurotransmitter serotonin, thereby making more serotonin available. Serotonin reuptake inhibitors have an inhibitory effect on the molecular transport system that removes serotonin from the synaptic cleft. This causes the absence of serotonin removal agents, which in turn leads to increased concentrations of serotonin in the brain. In contrast to the characterization of absences as being responsible for

mechanism failure, the absence of serotonin removal agents provides an example where the absence of an entity is causally responsible for *enhancing* the operation of a mechanism by making more serotonin available for postsynaptic binding. In line with the attenuator device analogy above, the absence in question increases the efficacy of transmission.

The point is that there are many ways in which an absence can have an effect on a mechanism. By appreciating this variety, we have a larger array of conceptual tools and a better understanding of causal efficacy to work with. The idea behind this section is “a weakened mechanism is still a mechanism” or conversely, “a strengthened mechanism is still a mechanism”.

To accord with empirical descriptions of mechanisms similar to those described above, we need a characterization of the effects of absence different to the one Barros provides. However, this does not mean that we should disregard his insights, especially the idea of mechanism failure and the use of baseline operations to compare facts. Instead, we need to expand on this thinking by including the understanding that the absence of an entity can both strengthen and weaken a mechanism by including instances of attenuation and enhancement in our explanations.

6. Functional variability

So far I have argued that the distinction between “function” and “failure” is inadequate because mechanistic phenomena are primarily situated in a continuum between these two extremes and, most often, not circumscribed (bound) by either. This observation, however, raises a deep question about mechanisms: At what point of attenuation does one mechanism become a different mechanism? Put differently, if mechanisms are defined and circumscribed by the phenomena they explain, at what point of attenuation does the phenomenon become different?

The function/failure distinction is encountered in several areas of philosophy of mechanisms. It is relied upon for differentiating between health and disease (e.g. Wakefield, 1992).⁶ It is appealed to for identifying functions in mechanisms (e.g. Craver, 2013). Among other things, the distinction is invoked for explaining causation by absence (Barros, 2013) and for separating mechanisms from non-mechanisms (e.g. Garson, 2013).

A body of recent literature has proposed that a mechanism serves a function by being a mechanism *for* something that contributes to the biological system in which it is contained. This means that a mechanism is circumscribed by what “it is supposed to do” (Craver, 2013, p. 138), which is achieved by “understanding how parts of organisms work, how they break or become diseased” (p. 154). Correspondingly, Stuart Glennan notes that “[w]hen one describes the behavior of a mechanism, one describes how it will behave if it is not broken” (2005, p. 448). And Justin Garson more recently adds, “heart disease is something that happens when this mechanism is disrupted” (2013, p. 319).⁷ The key point here is that thinking that a mechanism can behave differently or not at all follows from something that is *not* “constitutive of its functioning as such” (Craver, 2013, p. 9) reflects a line of reasoning that is ubiquitous in contemporary philosophy of mechanisms (see also, Bechtel & Richardson, 1993, p. 19; Darden, 2006, p. 259; Illari & Williamson, 2010, p. 285). Consequently, the idea of failure or malfunction has been central for explaining what mechanisms do when they are not doing what they are supposed to do. As Craver remarks, “much of physiological science such as neuroscience is driven not by the goal of understanding how the nervous system functions when it

⁶ According to Wakefield's analysis, disease/disorder is something that involves some kind of biological dysfunction, “where dysfunction means the *failure* to produce an effect selected during evolution” (Lemoine, 2013, p. 320, my emphasis).

⁷ Garson proceeds to suggest that “there is no mechanism for heart disease or Alzheimer's disease or schizophrenia because (for example) heart disease on part of a system is not a function of that system... since pathologies are almost universally held to be dysfunctional or non-functional states of a system” (pp. 320–322).

is working properly but rather by the goal of understanding how it can fail and how such failures might be predicted and controlled” (2013, p. 20).

Yet, reasons for the currently adopted scheme for circumscribing mechanisms are deeply problematic. Similarly to the previous issue of accounting for the variability in instances of functional decay in cases such as Parkinson’s disease, the source of this problem can be traced back to what is arguably an over-simplistic distinction between functional and non-functional/broken mechanisms.

So, for example, if heart disease is due to a disruption in a mechanism’s operation, then the question arises: By how much does it have to be disrupted before it develops into an abnormality? Considering that “normal” hearts pump 60–100 times per minute, it remains whether 101 beats per minute constitutes a pathological change. Similarly, should 59 beats per minute be interpreted as a malfunction? Indeed, providing normative descriptions of this kind has proven to be very elusive (McGrath, 2005). But it is especially difficult when presupposing two clearly delineated modes of functionality, which can be somewhat crudely characterized as on/off functioning. As indicated, there is no such thing as a sufficient amount of serotonin in the brain for it to function properly in the on/off sense of functionality. Rather, the continuum between sufficient and insufficient is best described as gradual.

The on/off approach to characterizing functionality has dominated the discussion to the extent that mechanism functionality appears to have two states: they are either functional or dysfunctional. However, as I have argued throughout, on/off accounts of mechanism function are implausible since they cannot account for degrees of breakdown and the gradual nature of dysfunction (i.e. mechanism attenuation). They also exclude a significant proportion of phenomena that cannot be appropriated within dichotomous formulations of functional variability. The function/failure dichotomy precludes the understanding of a crucial feature of mechanisms: attenuation or piecemeal variability in function. Insofar as mechanism failure is not an all or nothing feature, the idea of failure by degrees will not be compatible with the notion of causally productive mechanisms (as opposed to the idea of (absolute) failure). Although accounting for continuously varying phenomena might prove more challenging, it nevertheless offers a more accurate and fruitful approach to understanding biological and biomedical phenomena. Ultimately, mechanism attenuation presents many interesting line-drawing issues for contemporary philosophy of mechanisms.

7. Conclusion

The metaphysics of absences is a difficult issue. In this paper I suggest that a descriptively accurate account of scientific practice needs to acknowledge discussion about effects leading to vacillating forms of operational functionality.

Through the use of examples in this paper, I suggest that absences are causally responsible for a wide spectrum of effects on mechanistic operations, including attenuation and enhancement and failure to operate. More generally, these examples illustrate an important lesson from biological practice for accounts of causal explanation. We have seen that, at least when it comes to the biological sciences, one can speak of types of functionality, and degrees thereof, pertaining to the same biological mechanism affected by an absence.

I hope that the discussion in this paper has explained why Benjamin Barros’ idea of mechanism failure is not sufficient for a descriptively adequate account of causation by absence in the biological sciences. The effects of the absences in question are therefore best understood as part of the causal explanation for mechanism attenuation, enhancement and failure.

In conclusion, I think that the account developed in this paper captures a significant aspect of scientific practice as concerning most of the concepts scientists invoke to explain the effects of absences. More importantly, I hope this paper consequently shows that degrees of failure and mechanism attenuation are incompatible with the dichotomous distinction between failure/function that underwrites much of the contemporary philosophy of mechanisms.

Acknowledgments

I am grateful to John Cantwell and Jesper Jerkert for many valuable suggestions and discussions during the early stages of this work. Thanks are also due to Eva-Charlotta Mebius and two anonymous reviewers for their helpful corrections and much appreciated advice.

References

- Barros, D. B. (2013). Negative causation in causal and mechanistic explanation. *Synthese*, 190, 449–469.
- Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Science*, 36, 421–441.
- Bechtel, W., & Richardson, R. (1993). *Discovering complexity*. Princeton: Princeton University Press.
- Beebe, H. (2004). Causing and nothingness. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 291–308). Cambridge, MA: MIT Press.
- Bickle, J. (2006). Reducing mind to molecular pathways: Explicating the reductionism implicit in current cellular and molecular neuroscience. *Synthese*, 151, 411–434.
- Craver, C. F. (2006). When mechanistic models explain. *Synthese*, 153, 355–376.
- Craver, C. F. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. Oxford: Oxford University Press.
- Craver, C. F. (2013). Functions and mechanisms: A perspectivalist view. In P. Huneman (Ed.), *Functions: Selection and mechanisms* (pp. 133–158). Boston: Synthese Library.
- Darden, L. (2006). *Reasoning in biological discoveries*. Cambridge: Cambridge University Press.
- Dowe, P. (2004). Causes are physically connected to their effects: Why preventers and omissions are not causes. In C. Hitchcock (Ed.), *Contemporary debates in philosophy of science* (pp. 189–196). Malden: Blackwell.
- Evans, I., Thornton, H., Chalmers, I., & Glasziou, P. (2011). *Testing treatments. Better research for better healthcare* (2nd ed.). London: Pinter & Martin.
- Garson, J. (2013). The functional sense of mechanism. *Philosophy of Science*, 80, 317–333.
- Glennan, S. (1996). Mechanisms and the nature of causation. *Erkenntnis*, 44, 49–71.
- Glennan, S. (2002). Rethinking mechanistic explanation. *Philosophy of Science*, 69, S342–S353 (Proceedings).
- Glennan, S. (2005). Modeling mechanisms. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36, 443–464.
- Hall, N. (2004). Two concepts of causation. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 225–276). Cambridge, MA: MIT Press.
- Illari, P. M., & Williamson, J. (2010). Function and organization: comparing the mechanisms of protein synthesis and natural selection. *Stud Hist Philos Biol Biomed Sci*, 41(3), 279–291. <http://dx.doi.org/10.1016/j.shpsc.2010.07.001>.
- Illari, P. M., & Williamson, J. (2012). What is a mechanism? Thinking about mechanisms across the sciences. *European Journal of Philosophy of Science*, 2, 119–135.
- Kaplan, D. M. (2011). Explanation and description in computational neuroscience. *Synthese*, 183, 339–373.
- Lemoine, M. (2013). Defining disease beyond conceptual analysis: An analysis of conceptual analysis in philosophy of medicine. *Theoretical Medicine and Bioethics*, 34, 309–325.
- Locke, C., Fox, S., Caldwell, G., & Caldwell, K. (2008). Acetaminophen attenuates dopamine neuron degeneration in animal models of Parkinson’s disease. *Neuroscience Letters*, 439, 129–133.
- Machamer, P. (2004). Activities and causation: The metaphysics and epistemology of mechanisms. *International Studies in the Philosophy of Science*, 18, 27–39.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 60, 1–25.
- McGrath, S. (2005). Causation by omission: A dilemma. *Philosophical Studies*, 12, 125–148.
- Menzies, P. (2009). Platitudes and counterexamples. In H. Beebe, C. Hitchcock, & P. Menzies (Eds.), *The Oxford handbook of causation* (pp. 341–386). New York: Oxford University Press.
- Schaffer, J. (2004). Causes need not be physically connected to their effects: The case for negative causation. In C. Hitchcock (Ed.), *Contemporary debates in philosophy of science* (pp. 197–216). Malden, MA: Blackwell.
- Schaffer, J. (2005). Contrastive causation. *Physiological Reviews*, 144, 297–328.
- Torres, P. J. (2009). A modified conception of mechanisms. *Erkenntnis*, 71, 233–251.

Vesa, T. H., Marteau, P., & Korpela, R. (2000). Lactose intolerance. *Journal of the American College of Nutrition*, 19, 165S–175S.

Voss, M., Ingram, J. N., Haggard, P., & Wolpert, D. M. (2006). Sensorimotor attenuation by central motor command signals in the absence of movement. *Nature Neuroscience*, 9, 26–27.

Wakefield Jerome, C. (1992). The concept of mental disorder: On the boundaries between biological facts and social values. *American Psychologist*, 47(3), 373–388.

Woodward, J. (2011). Mechanisms revisited. *Synthese*, 183, 409–427.