Open-Ended Control vs. Closed-Ended Control: Limits of Mechanistic Explanation[1]

Jason Winning
Department of Philosophy, University of Toronto
jason.winning@gmail.com

> One of the beautiful things about mechanisms is, they don't think. They'll either function or malfunction.
>
> —Brad Bachelder, firearms expert, *Forensic Files* episode "Murder, She Wrote"

Abstract

Some recent discussions of mechanistic explanation have focused on control operations. But control is often associated with teleological or normative-sounding concepts like goals and set-points, prompting the question: Does an explanation that refers to parts or mechanisms "controlling" each other thereby fail to be mechanistic? In this paper I introduce and explain a distinction between what I call open-ended and closed-ended control operations. I then argue that explanations that enlist control operations to do explanatory work can count as mechanistic only if such control operations are closed-ended, not open-ended.

---

[1] This paper was delivered at the 2021 meeting of the Philosophy of Science Association.

1. Introduction

One of the virtues of mechanistic explanations has always been considered to be that they do not need to resort to a designer, goals, final causes, or other intentional or mentalistic-sounding notions. Mechanisms operate the way they do at any given moment solely because of the constituents they possess at that moment and how they are organized. According to the New Mechanistic philosophy of science, mechanistic explanations also refer to the "operations" or "activities" of the components and treat them as similarly describable in non-intentional, non-teleological terms.[2]

Recently, philosophers (e.g., Bechtel and Bich 2021) have called attention to the fact that mechanistic explanations in biology often include *control*: some of the operations of mechanistic components are control operations, and some mechanisms operate on others by controlling them. But control is often associated with teleological or normative concepts, such as that of a goal, set-point, or "normal range" of operation that the controller acts to achieve or maintain. Sometimes 'control' is even defined explicitly in terms of

---

[2] Some New Mechanists prefer to use the word 'entity' instead of 'component' or 'part', and 'activity' or 'interaction' instead of 'operation'. For purposes of this paper, I regard these terms as interchangeable. For discussion of this terminology, see Glennan (2017, 19–22).

paradigmatically intentional states like desires.[3] This is in contrast to other "stock in trade" (as Kauffman 1971, puts it) mechanistic operations like transport or ligand-induced conformational change, which are not thought of as being inherently normative or teleological. Can a mechanistic explanation that refers to parts or mechanisms "controlling" each other really be mechanistic?

My conclusion in this paper will be that whether or not control can factor into a mechanistic explanation depends on whether the control operation is understood in an *open-ended* or *closed-ended* way. An open-ended controller is one that does not operate with a fixed input–output relation but can be relied upon to select whatever the appropriate way to intervene on a target process is *because* of the appropriateness of that process. Mechanistic explanations can only include closed-ended control; the enlistment of open-ended control operations to do explanatory work renders an explanatory model as least partly teleological (and therefore at least partly non-mechanistic) in nature.

In what follows, I first introduce some reasons for thinking that it is important to make room for control operations in accounts of mechanistic explanation in biology. Next, I consider several arguments against inclusion of control in mechanistic explanations and show why they fail. Finally, I demonstrate the difference between understanding control operations

---

[3] For example: "A *control system* may be defined as a collection of interconnected components that can be made to achieve a desired response in the face of external disturbances" (Khoo 2018, 1).

in an open-ended and in a closed-ended way using an example and explain why mechanistic explanations should not include open-ended control operations.

2. Mechanistic Explanations and Control

"New Mechanist" philosophers of science such as Bechtel (Bechtel and Richardson 1993; Bechtel and Abrahamsen 2005), Glennan (1996, 2017), and Machamer, Darden, and Craver (2000) describe mechanistic explanations as explaining a system's production of a phenomenon of interest by reference to the causal operations that the system's parts perform (and that are performed on them), as well as the way these parts and operations are organized into a whole. This view of explanation was based on an analogy between biological mechanisms and human-built machines: "By calling the explanations *mechanistic*, we are highlighting the fact that they treat the systems as producing a certain behavior in a manner analogous to that of machines developed through human technology" (Bechtel and Richardson 1993, 17). While there is truth to this analogy, it can suggest a misleading view of biological mechanisms as mere *production* mechanisms whose construction and ongoing maintenance are taken for granted (Winning and Bechtel 2018). The internal operations of a human-built machine, an electric fan for example, are usually totally separate from the operations that went into initially constructing it or that go into fixing it if it needs repair. Its operation is "ballistic" in the sense that it carries out a predictable sequence of internal causal processes that are constrained merely by the device's own static internal structure; there is no need for internal *control*. Control of the system is usually done externally by a human operator; the device’s internal operation can be accounted for in a mechanistic way without reference to the notion of control.

Biological systems are very much unlike human-built machines in at least three ways:

- Whereas human-built machines resist thermodynamic equilibrium forces primarily due to the static rigidity of the materials out of which they are composed (metal alloys, plastics, etc.), organic systems are highly dynamic and maintenance of their structures far from equilibrium requires constantly varied work that must be internally controlled.
- Whereas human-built machines are built and repaired by external human agents, organic systems must continually construct and reconstruct their own structures, repair their own structures, and modify their own structures adaptively to maintain internal functioning in changing environmental circumstances.
- Whereas many functions of human-built machines are designed to function “ballistically” in the sense described above, internal control (especially feedback control) frequently occurs as one or more of the steps of an organic system’s internal production processes.

These features of biological systems are only possible due to the ubiquitous presence of control systems and control operations in biological mechanisms (Pattee 1973). But taking their role into account in New Mechanism is not a trivial matter of appending ‘control’ to the list of what kinds of operations a part in a mechanism might perform. The presence of internal control means that non-reducible, holistic *organization* takes on an even greater importance in mechanistic explanations, since the presence of feedback control loops introduces nonlinear phenomena that are difficult to explain or predict in terms of a static, sequential arrangement of parts and operations (Bechtel and Abrahamsen 2011). Further, control operations are, themselves, qualitatively distinct from other mechanistic operations and inherently more complex.

Given these facts, it might be wondered whether a new version of New Mechanism that gives a central role to control represents a modification to, or a departure from, mechanistic explanation properly so-called. In the following sections, I consider and evaluate several different arguments against inclusion of control in mechanistic explanations.

3. The "Black Box" Argument

One way to argue that control does not belong in mechanistic explanations derives from the inherent complexity of control operations. Components of mechanisms that perform control operations (i.e., controllers) have multiple parts that play various functional roles. When these are fully spelled out in the mechanistic model, they will be sufficient to explain what the controller does. Including both the controller itself as well as its subcomponents would then be redundant, and perhaps even a problematic form of overdetermination. To prevent this (the argument might run), mechanistic explanations should only include the subcomponents and their more basic operations, so that the controller and its control operations no longer feature in the model.

Of course, one can always get around this problem by omitting the mechanistic details of the controller and treating it as a single "black box"—a part within a larger mechanism that performs a control operation within that mechanism. However, it might be argued that a mechanistic explanation that does not detail the internal complexity of the black box would therefore be incomplete, or that this would prevent the explanation from being thoroughly mechanistic in nature. One might argue that control is what Craver (2006) refers to as a *filler term*, a term that is "used to indicate a kind of activity in a mechanism without providing any detail about how that activity is carried out" (2006, 360). According to Craver,

> filler terms are barriers to progress when they veil failures of understanding. If the term "encode" is used to stand for "some-process-we-know-not-what," and if the provisional status of that term is forgotten, then one has only an illusion of understanding. (2006, 360)

If the underlying mechanistic details of the controller are fully spelled out, then these details can simply replace the filler term 'control' in the explanation, so that it will no longer occur. To avoid having a mere mechanism "sketch" or mechanism "schema," and instead have a complete mechanistic explanation, the argument might run, one must therefore eliminate the control operations and replace them with more fully spelled out mechanistic details.

However, Glennan points out that 'complete' is a relative term when applied to mechanistic explanations. A mechanistic explanation need only be "complete at a single level of the mechanism" to avoid being a mere sketch or schema (Glennan 2017, 76). As long as the explanation details the operations at a given compositional level below where the phenomenon of interest occurs, such an explanation can be complete, even if the operations it specifies could themselves potentially be fleshed out further. A mechanistic explanation does not need to plumb the depths of every compositional level all the way down to elementary fields and particles. If control operations occur at a compositional level below the level of the whole that manifests a phenomenon of interest, the fact that such operations are themselves realized by sub-mechanisms does not mean that such operations cannot feature in a complete mechanistic explanation.

4. What Kinds of Operations Can Mechanistic Explanations Include?

But one might still argue that because control operations are more sophisticated than, and qualitatively distinct from, other kinds of mechanistic operations, they should not be included in a properly complete mechanistic explanation. Kauffman argued that operations appealed to in mechanistic operations should be "simple in the sense that it is by articulating together more than one of these well understood processes that we seek to explain more complex processes" (1971, 268). Correspondingly, it might be argued that referring to control operations does not sufficiently amount to explaining what is complex in terms of what is simple.

This raises the question of what kinds of operations are out-of-bounds in mechanistic explanations. While New Mechanist philosophers of science have sometimes listed examples of operations that might be included in a mechanistic explanation (e.g., "biosynthesis, transport, depolarization, insertion, storage, recycling, priming, diffusion, and modulation," Machamer, Darden, and Craver 2000, 8), they have generally avoided providing clear criteria for what can and cannot count as a part or operation included within an explanation without rendering the explanation non-mechanistic. One notable exception is the following passage:

> Though at times I adopt a specifically causal-mechanical view of explanation (see Craver 2007), and so will describe the ontic structures involved in explanation as causal or mechanistic, I intend the term *ontic structure* to be understood much more broadly. Other forms of ontic structure might include attractors, final causes, laws, norms, reasons, statistical relevance relations, symmetries, and transmissions of marks, to name a few. (Craver 2014, 29)

This suggests the following criterion: if an explanation relies on any of the things in Craver's list (attractors, final causes, laws, norms, reasons, statistical relevance relations, symmetries, or transmissions of marks) to carry explanatory heft, then that explanation is excluded from being entirely mechanistic. At best, it might be a hybrid of mechanistic and some other kind of explanation. We might also add that mechanistic explanations should not include things like historical properties or what Shoemaker (1980) calls "mere-Cambridge properties."

Important inclusions on Craver's list are final causes, norms, and reasons. More generally, explanations that rely on unreduced teleological or intentional properties to carry explanatory heft cannot be mechanistic. For example, suppose I am developing an alternative explanation of how the phenomenon of respiration occurs. I note that fumarase is clearly a key component playing a role in this process. On this basis, I decide that fumarase's presence must help to explain respiration because fumarase naturally does whatever would best facilitate respiration. Of course, this operation typically manifests in its reacting with fumaric acid, but I infer that if it could react with some other molecule that would be more helpful, it would do that instead. So in my explanation of respiration, I characterize the operation of fumarase as that it “does whatever is the best thing within its power to facilitate respiration.”

The problem here, of course, is that something like “do whatever is best” cannot count as a mechanistic operation (at least, not one that can do explanatory work in such an explanation), because it includes reference to what is “best,” a normative concept. Similarly, we could not characterize fumarase’s operation as being to “help secure the goal of respiration,” because this relies on the end state to help explain the occurrence of respiration, whereas mechanistic operations must refer only to states of the mechanism that occur during the production process, not after it has completed.

Suppose we instead say that the fumarase has the *current ability* (during the production process) to do whatever helps to secure the goal of respiration. This would at least avoid the appearance of backwards causation. The problem that remains is that the component, itself, is now being described as having a property verging on an open-ended general adaptiveness or instrumental rationality or practical wisdom such that it can be relied on to select the *appropriate* means of getting the job done, *because* it is the appropriate means. That a component can select a means *because* it is appropriate (rather than selecting the means that accidentally happens to be the appropriate one) is a form of guidance by goals or intentional states that cannot do explanatory work in a mechanistic explanation unless that explanation includes mechanistic details explaining how such guidance comes about. This is why Dennett (1973) argued that when behaviors are looked at from a purely mechanistic standpoint, they are necessarily treated as “tropistic,” or explainable without reference to appropriateness with respect to a goal.

It is important to note here that mechanistic explanations can include parts that have intentional states and the ability to select actions on the basis that they are appropriate for a goal, so long as these intentional states and abilities are not doing explanatory work in such explanations. For example, one can mechanistically explain how one group of people won a game of tug-of-war against another group of people by setting aside their goals and mental states and appealing simply to the physical forces exerted by the members of each group on the rope.

5. Mechanistic Accounts of Signal and Control Pathways

As philosophers like Bechtel and Bich (2021) emphasize, mechanistic explanations in biology often detail pathways of control. A control pathway involves detection of some condition, which triggers a signal pathway, which finally culminates in a change in gene expression or the operation of some particular mechanism which is the *target* of control. A fully detailed mechanistic description of a control pathway should explain what kinds of chemical reactions correspond to “detection” of the condition in question, the sequence of chemical reactions corresponding to the signal pathway, and the systematic relationship between detected conditions and the final control outcomes (e.g., changes in gene expression or change of the operation of some target mechanism(s)), referred to as the “input–output relation” (Shinar, Milo, Martinez, and Alon 2007). In order for this description to be mechanistic, the functional outcome of the signal pathway should be explained as a causal result of the signal pathway, and initial detection and its generation of the signal pathway should also be explained in causal terms, omitting teleological terms like ‘appropriate’, ‘too many’, or ‘enough’.

Given these restrictions, it might seem that no resources are left over to differentiate control from other mechanistic operations. Any mechanistic operation can be characterized as a “detection” in a certain sense, since it will be a response to the presence of some present condition or other. Any causal effect of such a “detection” can then be described as a “signal,” since it will carry information (in the sense of causal covariation) about the present condition that was detected. Any downstream causal effects of the “signal” might then be counted as “control” operations.

However, Bechtel and Bich argue that this would be to lose sight of the important difference between production and control. One way to distinguish control from non-control operations in mechanistic terms was suggested by MacKay, who wrote that in control systems,

> the input, A, determines the form of the output, B, without supplying all the energy of B … the energy of A is at least partly devoted to altering the structure through which the energy for B is channeled—altering the coupling between the output, B, and its internal energy supply (1964, 311)

The idea is that only if the signal pathway has a major qualitative effect on the mechanisms affected, while supplying a disproportionately low amount of the energy and/or resources necessary for that functional outcome, is it playing a *control* role.

An additional condition is necessary, however. Again, there must be a systematic relationship between the condition detected at the outset of the control process and the resulting qualitative effect on the controlled mechanism(s) (the input–output relation). A negative feedback control system, for example, has a qualitative effect that tends to affect the variable being measured so that it moves that variable closer to a particular value (the set-point). Many other systematic relationships are possible in control systems as well.

Importantly, though teleological considerations cannot enter into whether an operation in a mechanistic explanation counts as a control operation, and cannot contribute to the explanatory role of the operation, control operations and their functional outcomes can still be associated with teleological roles. For example, researchers might ask “How is the signal that the cell should divide generated?” and yet be asking for a mechanistic explanation for release of a certain signaling molecule, even though ‘should’ occurs in the question. A researcher might regard the signaling molecule as “teleological” in the sense of having the “purpose” of triggering cell division, but the characterization of it as having a purpose (rather than merely a causal effect) will carry no weight in the explanation. Such a mechanistic explanation does not actually require that the signal have, as semantic content, anything about what “should” happen.

Similarly, the set-point of a negative feedback control system might be characterized by a researcher as the "goal" of the system, but this attribution of a goal does not mean that the explanation of how the system works is thereby teleological in nature. A mechanistic explanation can treat the signal's tendency to bring the system towards the set-point merely as a causal disposition. The idea that the measured variable *should* be moved toward the set-point will not then play any role in explaining *how it does* move toward that point.

6. When Does Incorporation of Control Render a Model Non-Mechanistic?

Reliance on control operations can prevent an explanatory model from being mechanistic when those operations are understood as bestowing onto the controller an *open-ended* capacity to anticipate and execute whatever kind of modulation will be advantageous to the system. Control operations understood in this kind of way run afoul of Dennett's requirement that mechanistic operations must be "tropistic," referenced above. I now demonstrate what this would look like with an example.

Weber's Law is a systematic relationship that is observed in a wide range of kinds of visual systems of many species as they adapt to varying light levels without a corresponding decrease in resolution:

> As the background light level increases, the sensitivity of the visual system is decreased, which allows for operation over a huge range of light levels. From a dim starlit night to a bright sunny day, the background light level varies over 10 orders of magnitude (Hood and Finkelstein 1986), and yet our eyes continue to operate across

> all these levels without becoming saturated with light. The visual system accomplishes this by ensuring that its sensitivity varies approximately inversely with the background light, a relationship known as Weber's law (Keener and Sneyd 2009, 893)

This relationship is also observed in bacterial chemotaxis: "chemotactic cells … display 'logarithmic' tracking or sensing, characterized by a constant amplitude response when moving in a gradient that increases exponentially or nearly exponentially" (Sourjik and Wingreen 2012, 264). Suppose researchers of (fictional) lab X want to incorporate this characteristic into their mechanistic model of chemotaxis in *E. coli*. They identify an interaction network module of proteins that appears to be the location where control of the dynamic range of discrimination based on ambient level of ligand concentration likely occurs, and include this module in their model, though they are unable to determine the underlying details by which this control is executed. Suppose further that researchers from lab Y object to this mechanistic model on the grounds that merely including a "Weber's Law module" controlling dynamic range of discrimination is unilluminating; without detailing the underlying mechanism of that module, it adds nothing to our understanding of how chemotaxis works in *E. coli*.

The researchers of lab X respond by arguing that it is well known that Weber's Law represents an evolutionary optimization; we can infer that the module allows the system to maintain an optimal level of sensitivity to the changes in ligand concentration corresponding to the ambient level. Given an ambient level detection signal, we can predict the effect of the resulting sensitivity adjustment on flagellum behavior based on what would be optimal or most adaptive for the organism. The lab X researchers argue that these predictions are more robust given the module's inclusion as a part of the mechanism, rather than a mere inductive generalization from observed behaviors.

Of course, the researchers from lab Y would not find this argument satisfactory, and it is not an argument biologists would actually make. The problem is that the control operation itself is being defined or characterized in terms of what would be *optimal* or *most adaptive* for the organism, whereas of course a bacterium could not possibly have a control mechanism with information processing capabilities sufficient to entertain questions about what is optimal or most adaptive for it. Optimality models are a perfectly valid way of reasoning about biological systems, but the explanations they provide are teleological, not mechanistic, in nature.

We might characterize the operation of the module by saying that it "optimizes" the dynamic range of discrimination. Or we might characterize it by saying that it adjusts the dynamic range in a way that tends to correspond with Weber's Law. The first of these characterizes the control operation in teleological terms; the other characterizes it as a non-teleological, causal disposition. On the second characterization, the fact that the dynamic range of discrimination of changes in ligand concentration ends up being fairly close to what would be optimal for the *E. coli* is an accidental feature of the module itself, not an outcome it is part of the module's intrinsic feature set to produce. Even on the second characterization, however, inclusion of the module can add to the explanatory power of the mechanistic model in several ways. First, it localizes this functional role within the mechanistic model. Second, it places the operation within a sequential ordering of operations within the mechanism. Both of these allow for predictions about how the mechanism would respond to internal changes and for possibilities of experimental manipulation that mere knowledge that the dynamic range adjustment is observed to occur does not.

If a component in a biological system is found to play a control role that allows novel predictions, this is not necessarily because that component embodies an ability simply to do what would be optimal *because* it is optimal, or to modify some process in an adaptive way *because* it is adaptive. Control systems do not generally have anything approaching practical wisdom allowing them to just "know" the objectively best way to modulate whatever process they are modulating in given circumstances; no automatic tendency towards absolute optimality or absolute adaptiveness is implied by the fact that a component plays a control role in a system. It is only if something like this were implied by control that control would be problematic as a mechanistic operation.

7. Conclusions

Given the above considerations, we might define two ways that control operations can be invoked in a mechanistic explanation: closed-ended control versus open-ended control. Closed-ended control operates in a well-understood, predictable way, with a fully specified input–output relation (Shinar, Milo, Martinez, and Alon 2007). The input and output conditions are fully specified as a closed set of possibilities, as is the systematic relation that maps inputs to outputs. Open-ended control on the other hand, which is usually what is meant when discussing whether a human being is “in control” of something, means that the controller has a capacity to exert influence in a way that tracks *appropriateness* across an open-ended range of kinds of circumstances (in the language of philosophy of action, such control is “reasons-responsive”; McKenna 2017) and defies characterization in terms of a well-defined input–output relation.[4]

A key difference is that when treating a component as capable of open-ended control, you are making a normative claim: you are saying that the component, itself, has a kind of reliability that goes beyond causal regularity and is something like a rudimentary or proto-version of instrumental rationality or practical wisdom (without necessarily implying the presence of cognition or processing of representations). The component can be relied on (within some bounds) to select the appropriate means of getting the job done, *because* it is the appropriate means. The appropriateness of the means is then doing the explanatory work, not the fact that antecedent operations in the mechanism and the input–output relation of the

---

[4] A control system might incorporate some degree of stochasticity into its input–output relation, rendering its outputs probabilistic rather than deterministic. This could still count as closed-ended in my terminology if the probabilities are treated as fixed by non-normative factors.

controller were causally sufficient for it.[5] This (sub-cognitive analogy to) instrumental rationality or practical wisdom is taken for granted or even treated as brute or primitive, rather than treated as something calling for an explanation in mechanistic terms. The inclusion of such components introduces an unreduced *teleological* element in such explanations, and to that extent, renders such explanations non-mechanistic. The upshot is that as long as control operations are fully defined in a closed-ended way with a fixed input–output relation, their incorporation in an explanatory model is no obstacle to that model's being mechanistic.

References

Bechtel, William, and Adele Abrahamsen. 2005. "Explanation: A Mechanist Alternative." *Studies in History and Philosophy of Science Part C* 36 (2): 421–41.

Bechtel, William, and Adele Abrahamsen. 2011. "Complex Biological Mechanisms." In *Philosophy of Complex Systems*, ed. Cliff A. Hooker, 257–85. Amsterdam: North Holland.

Bechtel, William, and Leonardo Bich. 2021. "Grounding Cognition." *Philosophical Transactions of the Royal Society B: Biological Sciences* 376 (1820): 20190751.

---

[5] Some might of course argue that a regular input–output relation is, by itself, sufficient for a system embodying a norm. Even if true, this additional fact could not be employed to do explanatory work in a purely mechanistic model.

Bechtel, William, and Robert C. Richardson. 1993. *Discovering Complexity*. Princeton, NJ: Princeton University Press.

Craver, Carl F. 2006. "When Mechanistic Models Explain." *Synthese* 153 (3): 355–76.

Craver, Carl F. 2014. "The Ontic Account of Scientific Explanation." In *Explanation in the Special Sciences*, ed. Marie I. Kaiser, Oliver R. Scholz, Daniel Plenge, and Andreas Hüttemann, 27–52. Dordrecht: Springer.

Dennett, Daniel C. 1973. "Mechanism and Responsibility." In *Essays on Freedom of Action*, ed. Ted Honderich, 157–84. London: Routledge and Kegan Paul.

Glennan, Stuart S. 1996. "Mechanisms and the Nature of Causation." *Erkenntnis* 44 (1).

Glennan, Stuart S. 2017. *The New Mechanical Philosophy*. Oxford: Oxford University Press.

Kauffman, Stuart A. 1971. "Articulation of Parts Explanation in Biology and the Rational Search for Them." In *PSA 1970*, ed. Roger C. Buck and Robert S. Cohen, 257–72. Dordrecht: Reidel.

Keener, James, and James Sneyd. 2009. *Mathematical Physiology II: Systems Physiology*. 2nd ed. New York: Springer.

Khoo, Michael C. K. 2018. *Physiological Control Systems*. 2nd ed. Hoboken, NJ: John Wiley and Sons.

Machamer, Peter K., Lindley Darden, and Carl F. Craver. 2000. "Thinking About Mechanisms." *Philosophy of Science* 67 (1): 1–25.

MacKay, Donald M. 1964. "Cybernetics." In *Science in Its Context*, edited by John Brierley, 305–18. London: Heinemann.

McKenna, Michael. 2017. “Reasons-Responsive Theories of Freedom.” In *The Routledge Companion to Free Will*, ed. Kevin Timpe, Meghan Griffith, and Neil Levy, 27–40. New York: Routledge.

Pattee, Howard. H. 1973. “The Physical Basis and Origin of Hierarchical Control.” In *Hierarchy Theory*, ed. Howard H. Pattee, 71–108. New York: Braziller.

Shinar, G., R. Milo, M. R. Martinez, and Uri Alon. 2007. “Input Output Robustness in Simple Bacterial Signaling Systems.” *Proceedings of the National Academy of Sciences* 104 (50): 19931–35.

Shoemaker, Sydney. 1980. “Causality and Properties.” In *Time and Cause*, ed. Peter van Inwagen, 109–35. Dordrecht: Reidel.

Sourjik, Victor, and Ned S Wingreen. 2012. “Responding to Chemical Gradients: Bacterial Chemotaxis.” *Current Opinion in Cell Biology* 24 (2): 262–68.

Winning, Jason, and William Bechtel. 2018. “Rethinking Causality in Biological and Neural Mechanisms: Constraints and Control.” *Minds and Machines* 28 (2): 287–310.