

Mechanistic Explanation of Biological Processes

Abstract

Biological processes are often explained by identifying the underlying mechanisms that generate a phenomenon of interest. I characterize a basic account of mechanistic explanation and then present three challenges to this account, illustrated with examples from molecular biology. The basic mechanistic account is insufficient for explaining: 1) non-sequential and non-linear dynamic processes, 2) the inherently stochastic nature of many biological mechanisms, and 3) fails to give a proper framework for analyzing organization. I suggest that biological processes are best approached as a multi-dimensional gradient--with some processes being paradigmatic cases of mechanisms and some processes being marginal cases.

Contact Information

Philosophy Program, The Graduate Center, CUNY

365 Fifth Ave., Rm. 7113

New York, NY 10016

derek.skillings@gmail.com

Acknowledgements

I would like to thank Peter Godfrey-Smith, Massimo Pigliucci, and Arnon Levy for helpful comments and stimulating discussions. I am especially grateful to Shawn Simpson for providing illustrations. Thanks also to the many participants at the 2014 PSA who offered helpful feedback.

1. Introduction

The biological world is dynamic and evolving. One of its most prominent features is change, whether that change is an oscillating cell in homeostatic equilibrium, a developing multicellular organism, an evolving lineage, or seasonal ecosystem cycling. For much of biology providing a satisfactory explanation for a particular change means providing the *mechanism* that accounts for the phenomenon in question. A substantial amount of recent work in the philosophy of biology has been directed towards offering analyses of mechanisms and mechanistic explanation. Levy (2013) argues that this explosion of activity has led to mechanistic talk being adopted by three closely related, but often not well distinguished, projects (see Nicholson 2012 for a historical division of mechanistic views). The first project provides a mechanistic account of causation. It is the view that causal relations between non-fundamental physical phenomena exist in virtue of underlying mechanisms. The second project is concerned with explanation. A mechanistic view of explanation holds that certain types of phenomena, e.g. biological phenomena, are explained by identifying the mechanisms generating those phenomena, i.e. by specifying the underlying parts, their organization and their interactions. The third project is about how phenomena are best modeled. Levy (2013) uses the label “strategic mechanism” for the thesis that certain phenomena are most readily

understood when modeled abstractly as mechanisms. This paper is concerned with the mechanisms—and their underlying components—identified by the explanations in the second project, as well as the target systems to be modeled in a type three project. That is, I will primarily be focused on the entities underlying a phenomenon rather than the representation of those entities in a model of that phenomenon. I will leave to the side the discussion of the nature of causal relations, as it is, at best, tangential to my goals.

In this paper I will characterize a basic account of mechanistic explanation by presenting a recent and prominent approach in the philosophy of biology, almost universally referred to as the “new mechanist” approach (Section 2). I will highlight tensions between this basic approach and approaches that emphasize the non-sequential and stochastic character of some biological phenomena. Examples from recent research in molecular biology will be used to illustrate this point (Section 3). I will suggest that the best way to categorize the diversity seen across putative mechanisms is to view it as a multi-dimensional gradient—with some entities being clear cases of paradigm mechanisms and some entities being marginal cases—as opposed to being made up of discrete categories. I propose that this is accomplished by situating processes of interest within a multi-dimensional space generated by the interaction of three features: isolability, organization and sequentiality (Section 4). My overall goal is not to dispense entirely with the basic new mechanist approach,

but rather to find ways to extend the account so that it is adequate for handling the full diversity of biological processes that are putative mechanisms.

2. Mechanisms

The most widely cited philosophical account of mechanisms defines them 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, hereafter MDC). Other prominent accounts also define mechanisms by decomposing them in terms of parts and their activities or interactions: “...a mechanism is construed as generating a phenomenon through a start-to-finish sequence of qualitatively characterized operations performed by component parts” (Bechtel and Richardson 1993/2010; Bechtel 2011); “A mechanism for a behavior is a complex system that produces that behavior by the interaction of a number of parts, where the interactions between parts can be characterized by direct, invariant, change-relating generalizations” (Glennan 1996, 2002).

These three accounts serve as core examples of what is commonly called the *new mechanist* approach. Glennan focuses on parts, their properties, and their interactions. An interaction is just a change in the property of one or more parts that *brings about* a change in the property of another part. The ability to produce change is a feature of entities. MDC have an expanded ontology: entities don't just interact,

they engage in activities. Activities exist independently of entities. Activities are the producers of change, and entities are the things (with their properties) that engage in activities (MDC 2000). An important feature of activities is that they only exist extended in time, unlike entities, which can be extended in time but need not be. Bechtel and Richardson (1993/2010) emphasize how scientists decompose mechanisms structurally into their parts and functionally into their operations. They emphasize the same distinctions as MDC (entities/parts and activities/operations) but prefer the terminology of parts operating on each in order to highlight the direct involvement and multiple roles of parts, and are wary of making any metaphysical commitments. With these differences in mind, I will use the terms *basic mechanism* and *basic mechanistic account* to refer to a mechanism as outlined by the MDC account.

It will be helpful at this point to highlight with examples the features of a basic mechanism. The first feature of mechanisms is that they are hierarchically structured and can be broken down into components. The second important feature of a basic mechanism is that the entities in a mechanism must be appropriately located, structured, and oriented, and the activities in which they engage must have a temporal order, rate, and duration (MDC 2000). Activities connect entities into a coherent process. The organization of entities and activities determines the ways in which they are able to co-produce phenomena. For example, a neurotransmitter (entity) and a receptor (entity) bind (an activity) by virtue of their structural properties and charge

distributions. Similarly, a DNA base (entity) and a complementary DNA base (entity) hydrogen-bond (activity) because of their geometric structures and weak charges. Finally, the entities and activities constituting the mechanism must work in a regular fashion. A mechanism regularly produces or gives rise to the same phenomenon because earlier stages lead reliably to final stages, so that different instances of a mechanism share patterns of activity among similar or identical entities (Andersen 2012). What makes a mechanism regular is the productive continuity between stages (MDC 2000). Discontinuities are one way to mark the bounds of mechanisms.

An example of a basic mechanism comes from the text book description of the mechanism for protein synthesis in prokaryotes. At its simplest, the protein synthesis mechanism is a process in which many small entities, amino acids, are joined together by a ribosome to form a larger entity, a protein [figure 1 and the description that follows based on Moore (2012)]. The initiation of protein synthesis begins with the recruitment of ribosomal subunits from the cellular pool by initiation factors. The small ribosomal subunit then binds to an appropriately aligned mRNA and an fMet tRNA. Initiation ends when the second amino acid in the protein sequence is carried to the ribosomal complex by a tRNA. Elongation of the polypeptide then ensues in a cyclical fashion. During elongation, a tRNA carrying an amino acid binds to the ribosome in what is called the A site. The ribosomal complex moves the tRNA from the A site into the P site, and another tRNA carrying the $n+1$ amino acid moves into the A site. Within the P site the ribosome then transfers the amino acid from the

tRNA onto the end of the growing protein in a process called peptidyl transfer. At this point the ribosome must be reinitialized before the next amino acid can be added to the nascent protein. During this process, which is called translocation, the amino acid-less tRNA in the P site leaves the ribosome, the $n+1$ tRNA in the A site moves to the P site, and the ribosome moves along its mRNA in the 3' direction by one codon. When the translating ribosome encounters a stop codon in the mRNA, the protein is released from the ribosome and the two ribosomal subunits return to the cellular pool. This process is called termination.

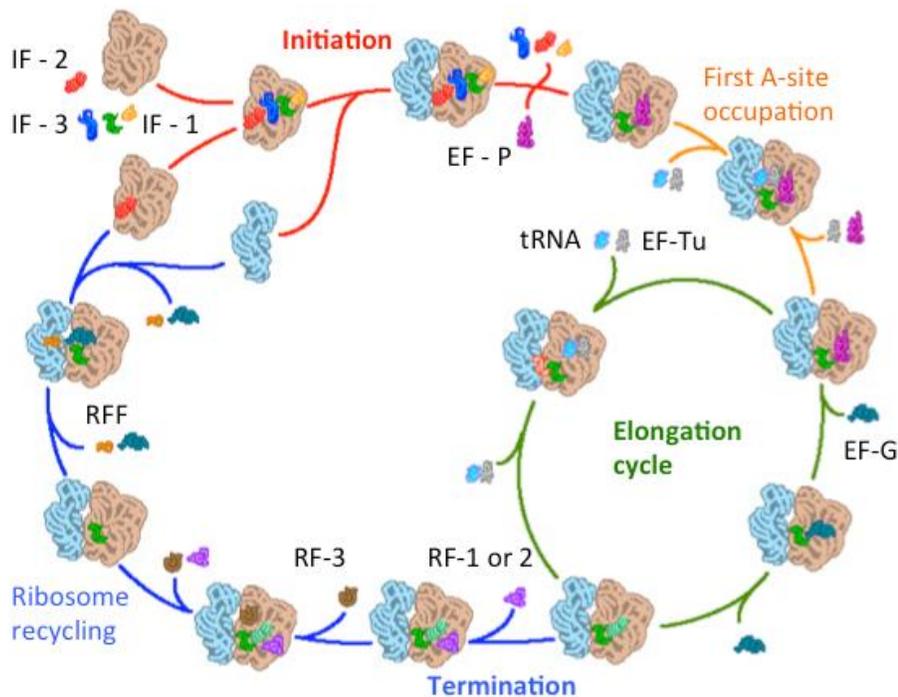


Figure 1. Protein synthesis cycle based on Moore (2012). Illustration by Shawn Simpson.

3. Criticisms of the Basic Mechanistic Account

Biological explanations that adhere to a basic mechanistic account like the one above have often been criticized as inadequately representing and explaining many biological phenomena. I will group these criticisms into three types. First, treating mechanisms as paradigmatically sequential and linear fails to accommodate a vast number of dynamic biological processes that are non-sequential, oscillatory, branching or feedback systems (Bechtel and Abrahamsen 2005, 2013; Bechtel 2011). Second, a strong requirement for regularity in mechanisms fails to accommodate the biological processes that are internally and temporally stochastic as a result of their dependence on the levels of free energy in the system. Finally, the basic mechanistic accounts fail to adequately distinguish between processes that are organized in an orderly way with parts that have differential roles and local relations, and those processes in which such features are absent (Levy and Bechtel 2013, Levy 2014).

3.1. Non-Sequential Dynamic Mechanisms

The basic mechanistic account emphasizes that a mechanism moves from start or set-up conditions to finish or termination conditions in an ordered and sequential way (MDC 2000). In a series of papers, William Bechtel and Adele Abrahamsen argue

that the basic mechanistic account is limited by its emphasis on sequences of operations (Bechtel 2011, 2012, 2013; Bechtel and Abrahamsen 2005, 2013). They advance the term *dynamic mechanistic explanation* for an extended account that includes non-sequentially organized mechanisms.

The most straightforward and tractable organization that a process can have is one in which there is spatial adjacency and connectedness in its parts, and a stepwise temporal sequence in its operations. Machine-like mechanisms are often organized in this way. More complex sequential organizations are also common. The Krebs cycle is not a start to finish sequence, but rather a sequence that runs in a loop. Many biochemical mechanisms are also built around some sort of negative feedback loop, a process in which the result of the process regulates the process itself.

Understanding these more complex sequences requires one to become increasingly more concerned with the dynamics of mechanisms, or how changes in the environment can affect the rate, duration and activation of a mechanism. The basic mechanistic account can accommodate this low-level dynamical behavior as long as it sequential. But many biological mechanisms exhibit oscillations or more complex dynamical behavior, and this can be crucial for orchestrating operations within the mechanism (Bechtel and Abrahamsen 2013). A challenge to the basic mechanistic account comes from complex dynamical behavior that emerges when mechanism organization is *non-sequential*. Organization might be non-sequential because there are forks that can take the mechanism in different directions. Non-

sequential organization might also arise from redundant but different sub-mechanisms, such that final product or state of the primary mechanism is unaffected, but the path to that point can differ. Given a particular end-state it might be impossible to determine the specific path a non-sequential mechanism traveled to get there. This is not the case for a sequential mechanism. This doesn't mean that a non-sequential mechanism doesn't produce regular changes, or that the parts and relations of the mechanism are indescribable. Rather, it is a point about how in any given instance of the mechanism, the unfolding activity of that mechanism might differ in response to differences in start-up conditions or ongoing environmental inputs.

3.2. Stochastic Mechanisms

Under the basic mechanistic account, a process must behave regularly in order for it to be considered a mechanism. What makes a mechanism regular is the productive continuity between stages (MDC 2000). A mechanism regularly produces or gives rise to the same phenomenon because earlier stages lead reliably to final stages, so that different instances of a mechanism share patterns of activity among similar or identical entities (Andersen 2012). A strong requirement for regularity in mechanisms fails to accommodate the biological processes that are internally and temporally stochastic, such as when their mode of operation is highly dependent on the levels of free energy in the system. In these types of systems, changes are

probabilistic rather than regular. In many cases, instances of a stochastic mechanism do not share the exact same pattern of activity among even identical entities.

Let's return to the example of protein synthesis. The mechanism outlined in Section 1 was a version of the classic textbook explanation of protein synthesis. The ribosome is one of the primary players in the mechanism of protein synthesis. As more has been learned about it, the picture of protein synthesis has begun to change. Moore (2012) characterizes the progress in describing the ribosome since the first atomic resolution crystal structures of ribosomal subunits were published in 2000. The explicit goal of many structural biochemists was and is to provide a "movie" of protein synthesis. In this project the ribosome is characterized at discrete stages. The protein synthesis movie is created by characterizing all the different states of the ribosome and how they are related to each other, these are still frames of the movie that are put in order and played to show the process of protein synthesis. At the heart of the movie is a ribosome that is doing something. There are regular changes between stages, and earlier stages lead reliably to later stages.

Moore argues that these structure-based movies of the ribosome make it appear to be something that it is not. The reason that the movies are misleading is because ribosomes don't behave like machines that can be explained mechanically. Moore states that: "[A]ll the functionally significant movements of the ribosome, both internal and external, are biased random walks, and it is most unlikely that any given ribosome will ever do exactly the same thing twice as it elongates some polypeptide"

and “It follows that if a movie of elongation is not to be totally misleading, it must depict the endless series of meaningless, thermally driven, conformational fluctuations that separate one functionally significant event from the next.” He predicts that 99.99% of the activity in protein elongation is random fluctuation that has nothing to do with function and that most of the progress through the protein elongation cycle is driven by changes in relative free energies coming from conformational state changes.

Think of a machine like a mechanical clock. The workings of the mechanical clock are explained by looking at how the physical parts of its mechanism interact. The mainspring stores mechanical energy through torsion. As the mainspring unwinds it rotates a central gear, which in turn drives the rest of the gear train. The last gear engages the escapement. The escapement is also engaged by the balance wheel, a weighted wheel which keeps time by rotating back and forth at fixed intervals according to the tuning of a balance spring which pushes back on it. The interaction between the mainspring and the balance wheel via the escapement mechanism turn the gears at a fixed rate. These gears are attached to the rotating hands of the clock face, which indicate the time. The movement and interactions of the parts of the watch explain how the watch works.

The parts of a protein, like a ribosome, don't stand in the same relations as the parts of a mechanical clock. The mainspring of a watch can reliably turn the gear train. Macromolecules immersed in liquids are at the mercy of frictional and

thermodynamic forces. At any given time, the forces produced by a ribosome are one-tenth of the level or less than the average thermodynamic forces experienced by that ribosome. This means that the ribosome is randomly pushed into different conformational states by external forces, rather than moving in sequence according to the activity and interactions of its parts. Thus, all the functionally significant activities of the ribosome, both internal and external, in protein synthesis are probabilistic (Moore 2012). The probability that a ribosome will engage in an activity that is important for protein synthesis depends on the free energy in the system and the presence of protein factors that promote some states over others.

Moore recommends that we “come to think of the ribosome as the dynamic, constantly varying structure it is and always was” and move away from quasi-mechanical explanations for the properties of ribosomes and “instead seek to understand them in terms of the particle’s conformational energy landscape.” He boldly claims that structure-based movies of macromolecular processes have no explanatory power.

The mechanism of protein synthesis doesn’t neatly fit into the basic mechanistic account. There is no productive continuity between stages, where earlier stages directly produce later stages. Rather, the mechanism proceeds stochastically. Movement through the process of protein synthesis is a biased random walk. Watches reliably tell time, and ribosomes reliably produce proteins, but the relations between the temporal stages of the mechanisms are much different. A watch has a fixed

structure and operates in a fixed sequence. The parts of the ribosome are constantly in motion and changing their relations to each other due to forces that dominate at the molecular level, such as: atomic bonds, friction, thermal forces and internal vibrational decay. These ongoing interactions are what dynamically stabilize the ribosome even as it flows constantly from conformation to conformation in a biased random walk. Furthermore, it is this constant change that leads to the functionally important states needed for protein synthesis. The ribosome is a dynamic, constantly varying entity, and it is this feature that is most explanatory when it comes to understanding how the ribosome synthesizes proteins.

3.3. Order in Mechanisms

The basic mechanistic account requires that the entities in a process are organized if they are to be considered a mechanism: parts must be appropriately located, structured, and oriented, and the activities in which they engage must have a temporal order, rate, and duration. This account fails to adequately distinguish between processes that are organized in an orderly way with parts that have differential roles and local relations, and those processes in which such features are absent (Levy and Bechtel 2013, Levy 2014).

Consider two processes in the mechanism of Ca^{2+} regulation: the passive diffusion of Ca^{2+} from the interior of the mitochondria out into the cytosol, and the

active channeling of Ca^{2+} across a membrane. In the case of diffusion there is a system that consists of a semi-permeable membrane that straddles an electrochemical gradient. The exterior environment is more negatively charged than the interior environment and the positive Ca^{2+} ions will move across the membrane until electrostatic equilibrium is reached. The rate of diffusion is primarily a factor of the permeability of the membrane and the difference in charge between the two environments.

In the case of active channeling, when Ca^{2+} levels in the cytoplasm drop, secondary messaging systems within the cell can activate protein complexes that release Ca^{2+} stores from inside the endoplasmic reticulum. This example describes ryanodine receptor 2 (RYR2). RYR2 is a membrane-bound complex composed of four subunits that form a channel for Ca^{2+} to cross (figure 2). The complex is associated with various proteins that function to modulate its opening including: factors (like CSQ) that modulate the sensitivity of RYR2, transmembrane proteins that facilitate the opening, multiple enzymes that bind and control the “powering” of RYR2 by cAMP, and scaffolding proteins that bind further enzymes that reset and close the channel.

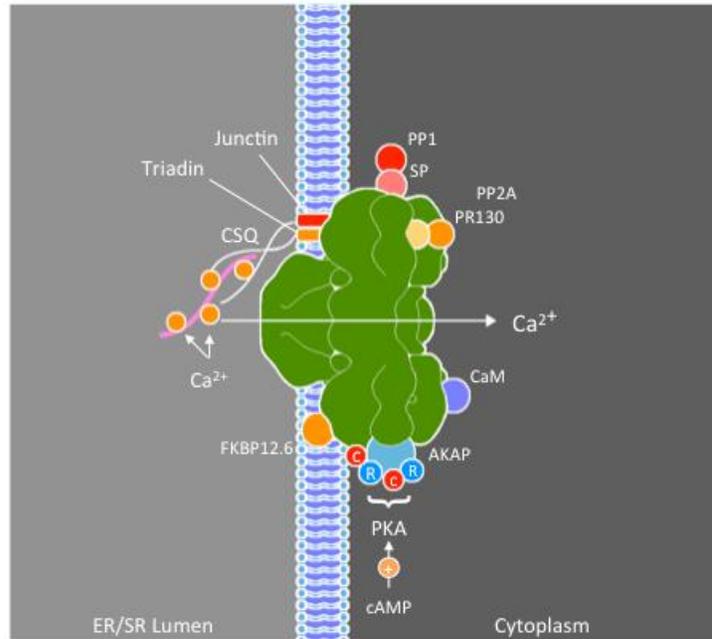


Figure 2. Ryanodine receptor 2 (RyR2) based on Berridge et al. (2003). Illustration by Shawn Simpson

Both of these example processes are dependent on system-level organization and feature interactions between constituent parts. One important difference is in how they are ordered. Active channeling is a highly ordered phenomenon: every component must be in the right place, at the right time, playing the right role. Diffusion, on the other hand, is a disorderly phenomenon: the role and layout of individual particles is insignificant. It doesn't matter what ions help generate the electrochemical difference across the membrane from the high Ca^{2+} concentration, nor does it matter what the composition of the membrane is as long as it is semi-

permeable. All that matters are the average properties of the components of the system.

Levy (2014) gives a characterization of orderliness that I will adopt here.

Suppose we have a system *S*, exhibiting a behavior *B*. *S* is orderly to the extent that:

- (a) Distinct components of *S* play different roles in bringing about *B*.
- (b) Components play their roles in virtue of local relations to other components.

An orderly system exhibits an internal division of labor, analogous to that present in many manmade machines: each part does something distinct and recognizable, but there is also interdependence among parts, so that the system's overall behavior is an integrated product of their activities. Component parts can be distinct in two ways.

They might be distinct because they are spatially separated. They might also be functionally distinct, i.e., they are distinct with respect to the difference made by the components. Levy fleshes this out in terms of Woodward's conception of modularity (Levy 2014, Woodward 2003). A component's contribution is modular if the difference it makes is independent of the difference made by other components. That is: component *X*'s contribution is modular if it is possible to disrupt the activity of other components without affecting the contribution of *X*. In RYR2 channeling, for instance, it may be possible to deform the transmembrane proteins that help stabilize the size of the channel opening, thus disrupting channel size, without affecting the components that flip the RYR2 complex between open and closed states, and vice

versa. These would then be distinct components of the active channeling mechanism. This functional distinctness in roles is what contributes to higher levels of orderliness.

The requirement that components play their roles in virtue of local relations to other components is related to the idea that an orderly system exhibits internal integration. It is possible for a system to have functionally distinct parts, i.e. for its parts to contribute differentially, without those parts being integrated. For example, the electrochemical charge of the cell is determined by the mixture of positive and negative ions within the cytoplasm. There are many different components that play different roles, but they do so by freely mixing, such that the specific layout of elements does not matter much. It is an aggregate property of the system (Wimsatt 1986). This is different than an active membrane channel system, there the local relations between parts are very important for the proper functioning of the process. In regards to electrochemical charge, it won't matter much if you switch two Ca^{2+} ions on opposite sides of cell with each other, or if you swap a Ca^{2+} for two H^+ or a Mg^{2+} . In contrast, an active membrane channel will be greatly disrupted if a transmembrane protein is swapped with a receptor that binds cAMP.

4. Dealing with Mechanistic Diversity

I have given a basic account of mechanistic diversity and then categorized three ways in which this basic account can be found to be lacking. None of the

criticisms I have discussed have challenged the existence of mechanistic phenomena, but rather find the basic account inadequate for handling the diversity of processes that are putative mechanisms. One route to dealing with this diversity of mechanism-like phenomena is to try to and isolate *the* true mechanisms by giving a more substantial theory. Another route is to subdivide mechanisms into smaller discrete categories: machine mechanisms, biological mechanisms, dynamic mechanisms, etc. A third approach begins by recognizing that there is gradient from phenomena that are produced by processes that are paradigmatically machine-like in their mechanistic functioning, to phenomena that are produced by processes that don't conform to even the most minimal and basic account of mechanisms. The approach here is not to try and divide this space up discretely, but to describe the dimensions along which these phenomena grade into each other. I offer a brief account of mechanisms in the spirit of this third approach following the example set by the treatment of biological populations in Godfrey-Smith (2009).

Once again, my starting point is the basic mechanistic account. The minimal goal in mechanistic explanation is to explain a phenomenon by identifying the lower-level phenomena that produce the higher-level target phenomena. This requires identifying the lower-level entities and characterizing what they do and how they interact with other. This is the minimal notion I will use for what is required to consider something as even a putative mechanism: a hierarchically structured system where the higher-level can be at least partially decomposed into lower-level parts.

Not all phenomena with these characteristics will be mechanisms, but it is possible to give the features of systems which distinguish the clear or paradigm cases of mechanisms from more marginal ones. Three interacting parameters will be used to make this distinction (figure 3).

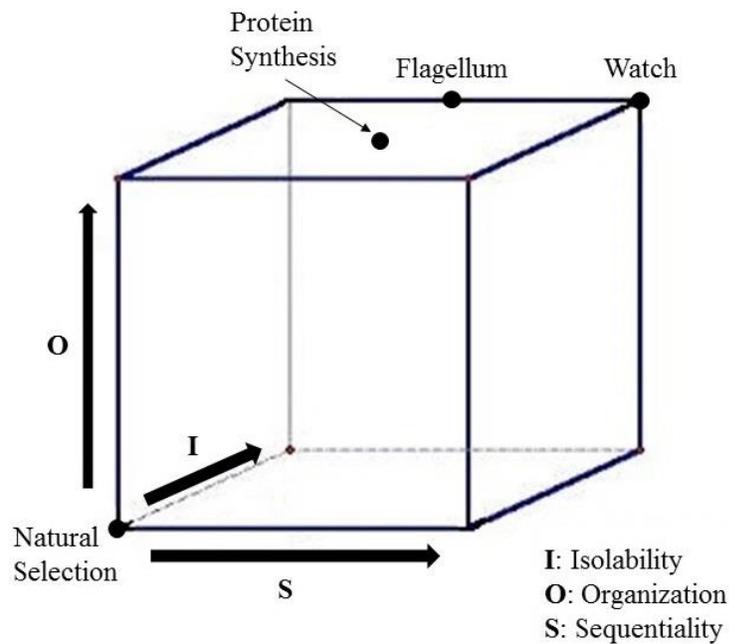


Figure 3. Three dimensions of mechanism-related processes.

The first parameter is the isolability of the system. This is the measure of the overall sharpness of the boundary between the system and its environment.

Identifying the start-up, maintenance and finishing conditions of a process requires being able to identify what is part of the process and what is input from, or output to, the environment. An example of a system with sharp boundaries is a mechanical

watch. An example of a system with vague boundaries is a population undergoing speciation with geneflow. An intermediate case might be synapse firing.

The second parameter is the organization of the system. This is the measure of the orderliness of the system in terms of the distinctiveness of its component parts, the discreteness of its component parts, and the importance of local relations between parts in producing the higher-level phenomena. Highly ordered systems include bacterial flagella, and again, a mechanical watch. A system with low order is two liquids undergoing passive diffusion across a membrane, and a system with intermediate order might be a population undergoing selection.

The third parameter is the sequentiality of the system. This is a measure of the linearity and stochasticity of the system. A system with low sequentiality might be characterized by feedback loops, cycling, oscillations, variation in how component parts interact, or stochastic activities. A system with low sequentiality is the process of Ca^{2+} homeostasis. A moderately sequential system is the process of protein synthesis. A highly sequential system is a coin-operated gumball machine.

The clearest cases of mechanisms are machine mechanisms like a mechanical watch, which have a high degree of isolability, organization and sequentiality. These are the systems that most clearly meet the criteria of the basic mechanistic account. An example that is at best a marginal case of a mechanism is evolution by natural selection, a process that has a lower degree of isolability, organization and sequentiality. Many biological processes that are taken to be true mechanisms are

highly isolable and ordered, but are only moderately sequential. There are a large number of intermediate cases.

There is an increasing recognition of the importance of cyclical, oscillating and biased stochastic processes for the explanation of biological phenomena. Using this multi-dimensional framework could help to extend the basic account of mechanistic explanation so that it can deal with biological mechanisms that are non-sequential, stochastic or less-orderly.

References

Andersen, Holly. 2012. "The Case for Regularity in Mechanistic Causal Explanation." *Synthese* 189:415-432.

Bechtel, William. 2011. "Mechanism and Biological Explanation" *Philosophy of Science* 78(4):533-557.

Bechtel, William. 2012. "Understanding Endogenously Active Mechanisms: A Scientific and Philosophical Challenge." *European Journal for Philosophy of Science* 2(2):233-248.

Bechtel, William, and Adele Abrahamsen. 2005. "Explanation: A Mechanist Alternative." In: *Studies in the History and Philosophy of the Biological and Biomedical Sciences* 36:421-441.

Bechtel, William, and Adele Abrahamsen. 2013. "Thinking Dynamically About Biological Mechanisms: Networks of Coupled Oscillators." *Foundations of Science* 18:707-723.

Bechtel, William, and Robert C. Richardson. 1993/2010. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Repr., Cambridge, MA: MIT Press.

Berridge, Michael J., Martin D. Bootman, and H. Llewelyn Roderick. 2003. "Calcium Signalling: Dynamics, Homeostasis and Remodelling." *Nature Reviews: Molecular Cell Biology* 4:517-529.

Glennan, Stuart. 1996. "Mechanisms and the Nature of Causation." *Erkenntnis* 44:50-71.

Glennan, Stuart. 2002. "Rethinking Mechanistic Explanation." *Philosophy of Science* 69:S342-S353.

Godfrey-Smith, Peter. 2009. *Darwinian Populations and Natural Selection*. Oxford: Oxford University Press.

Levy, Arnon. 2013. "Three Kinds of New Mechanism." *Biology and Philosophy* 28:99-114.

Levy, Arnon. 2014. "Machine-Likeness and Explanation by Decomposition." *Philosopher's Imprint*. 14(6).

Levy, Arnon, and William Bechtel. 2013. "Abstraction and the Organization of Mechanisms." *Philosophy of Science* 80(2):241-261.

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

Moore, Peter B. 2012. "How Should We Think About the Ribosome?" *Annual Review of Biophysics* 41:1-19.

Nicholson, Daniel J. 2012. "The Concept of Mechanism in Biology." *Studies in History and Philosophy of Biological and Biomedical Sciences* 43:152-163.

Wimsatt, William C. 1986. "Forms of Aggregativity." In: *Human Nature and Natural Knowledge*. (Eds.) Alan Donagan, Anthony N. Perovich, and Michael V. Wedin, 259–93. Dordrecht: Reidel.

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.