

Understanding Biological Mechanisms: Using Illustrations from Circadian Rhythm Research

William Bechtel

1 Introduction

In many fields of biology, when investigators seek an explanation for a phenomenon, what they are seeking is an account of the mechanism responsible for it. The search for mechanisms to explain phenomena has played an important role in biology for over two centuries. Twentieth century philosophy of science, however, largely neglected mechanisms; the dominant account of explanation held that explanation involves deriving descriptions of phenomena from statements of laws and initial conditions (Hempel 1965). Noting that laws are seldom averted to in biological explanations (see Lange this volume for laws in biology) but references to mechanisms are ubiquitous, in the last two decades several philosophers of biology have turned their attention to characterizing what biologists take mechanisms to be and the strategies they employ to discover, represent, and evaluate mechanistic explanations (Bechtel and Abrahamsen 2005; Bechtel and Richardson 1993/2010; Machamer et al. 2000). These philosophical analyses of mechanistic explanation can be valuable for educators seeking to present the framework of biological inquiry to students (van Mil et al. 2013).

The distinguishing feature of mechanistic explanation is that scientists seek to explain a phenomenon of interest by identifying the working parts of the responsible mechanism—the parts that perform the various operations that go into producing the phenomenon. Since mechanistic investigation proceeds by decomposing a mechanism into its component parts, which in a straightforward sense are at a lower level of organization (they are necessarily smaller than the mechanism as a whole), it is often characterized as reductionistic. This sense of reduction is rather different

W. Bechtel (✉)

Department of Philosophy and Center for Chronobiology, University of California,
9500 Gilman Drive, La Jolla, San Diego, CA, USA
e-mail: bill@mechanism.ucsd.edu

from that often presented in the philosophical literature (Nagel 1961) in which laws characterizing phenomena at one level are derived from those at a lower level and according to which in the end everything might eventually be explained from the lowest-level. For one thing, the parts and operations appealed to in mechanistic explanations may not be characterized in terms of laws. But even more fundamentally, the lower level does not provide all the information needed to account for the phenomenon. The working parts only produce the phenomenon when they are organized and their operations are appropriately orchestrated. Many components of biological mechanisms operate differently when situated within the mechanism than when removed from it (Boogerd et al. 2005).

Knowledge of the manner in which a mechanism is organized and how it affects the parts is additional to the knowledge of the parts that is gleaned from focusing on them treated individually. Moreover, the conditions under which the mechanism works may depend on conditions imposed on the mechanism from its environment. Thus, the reductionistic knowledge provided by mechanistic decomposition requires additional knowledge of the modes of organization realized at higher levels, including levels in which the whole mechanism is just a part, in order to explain a phenomenon. Accordingly, although the techniques for doing so are less developed than those for decomposing mechanisms, mechanistically oriented scientists must also recompose and situate mechanisms in order to account for how they produce the phenomenon. Drawing students' attention to both the valuable aspects of reductionistic decomposition and the need to complement decomposing with recomposing and situating mechanisms can help them develop a more comprehensive understanding of biology (see also Braillard this volume for reductionism and systems biology).

After first further articulating the nature of mechanistic explanation in the next section, I will in subsequent sections discuss the key tasks in developing such explanations—delineating the phenomenon, identifying and decomposing the responsible mechanism, and recomposing and situating the mechanism. To help make the exposition concrete, I will develop as an example throughout the chapter research on circadian rhythms—endogenously controlled oscillations of approximately 24 h in many physiological processes (e.g., basic metabolism and body temperature) and behaviors (e.g., locomotion and cognitive performance). Circadian rhythms are a fascinating phenomenon that readily attracts student interest and research on them provides a reasonably accessible case for introducing students to the intricacies of mechanistic research.

Since circadian rhythms are maintained endogenously, researchers have sought an internal mechanism, a clock, to explain them. Because rhythms are manifest even in single-celled organisms (cyanobacteria and fungi), researchers have assumed that the clock mechanism exists inside individual cells and over the last 20 years an explanatory schema involving a transcription-translation feedback loop (in which a product of gene expression builds up and as it does so inhibits gene expression until it is degraded) has received strong support. In cyanobacteria Nakajima et al. (2005) demonstrated that a system consisting of the KaiA, KaiB, and KaiC proteins together with ATP sufficed to generate circadian oscillations, pointing to a simpler feedback loop involving phosphorylation and dephosphorylation of a crucial protein as sufficient for

circadian rhythms. But even in cyanobacteria, the transcription-translation loop is assumed to play a fundamental role in normal maintenance of circadian rhythms. In multi-celled organisms the components of the clock mechanism are found and exhibit the appropriate behavior in (nearly) all cells of the organism, but in animals the ability to maintain circadian rhythmicity depends upon the clock mechanism in specific cells in the brain. In mammals, these cells are located in a small structure within the hypothalamus known as the suprachiasmatic nucleus (SCN). As I will discuss below in the context of discussing the process of recomposing a mechanism, there is a complex coordination process whereby individual SCN neurons depend on others to maintain a reliable circadian rhythms.

2 Characterizations of Mechanisms

The ideas scientists deploy to explain nature often originate as metaphors in which they transfer a framework humans have developed in technological applications to understand a system they encounter in nature. The recent attraction of the computer metaphor for understanding the mind and brain is an illuminating example. Once humans had constructed machines that could manipulate symbols (encoded as strings of 0s and 1s), cognitive scientists and neuroscientists began to apply the idea of computation to characterize how information is processed. (Over time, the idea of computation has been extended beyond that which Turing (1936) and Post (1936) initially proposed, and now often seems to involve nothing more than a series of causal processes that transform one representation into another.) This example illustrates a more general practice in which scientists have tried to understand natural processes using ideas developed in the context of humanly constructed machines.¹ Descartes was one of the first to give clear expression to this idea, characterizing

¹Pigliucci and Boudry (2011) argue that what they call the machine-information family of metaphors has negative consequences in both science education and scientific research (see also Brigandt this volume). In particular, in science education it provides an opening for intelligent design. While the comparison to human made machines does invite appeals to a designer, rejecting the practice of several hundred years in biology of seeking mechanisms to explain phenomena is not a good educational strategy. A better strategy is to consider seriously the sort of mechanisms that are found in living organisms (e.g., ones that build and repair themselves) as well as to emphasize that machines are typically not optimally designed, and this is especially true of biological mechanisms. The origin of biological mechanisms is better explained by evolutionary processes (drift as well as selection) than by appeal to an omniscient intelligent designer. A further part of Pigliucci and Boudry's critique focuses on the use of the information or blue-print metaphor for the relation of genes to biological traits. They appeal to work focusing on biological development to show that the relation between genes and traits is far more complex—organisms and environments figure centrally in explaining how genes are expressed. This is the view of developmental systems theory and one way proponents such as Griffiths and Gray (1994) present their message is by viewing genes as just one part of the complex developmental mechanism responsible for the appearance of traits in organisms in successive generations.

“this Earth and indeed the whole universe as if it were a machine.” For Descartes a machine operated as a result of its parts having specific shapes that affected each other’s movement brought about by contact action, and so he continued: “I have considered only the various shapes and movements of its parts” (*Principia* IV, p. 188). Descartes extended this idea to the bodies of organisms, proposing to explain all animal functions and behavior (including the same behaviors when they occurred in humans) in terms of the push and pull of component parts. Many investigators followed in Descartes’ footsteps, but unlike Descartes, engaged in empirical investigation as they sought to understand biological mechanisms. Moreover, over time they extended the range of basic processes they viewed as appropriate for introducing in mechanistic explanations to include Newtonian attraction between objects, the creation and breaking of chemical bonds, and conduction of electrical currents.

The philosophers of science who have begun to focus on mechanistic explanation have been concerned primarily to provide an account of explanation that fits the practice of biologists. Thus, they have often started with particular examples of explanations found in biology: explanations of basic metabolic processes such as fermentation, of the synthesis of proteins, or of the transmission of chemical signals between neurons. They have then noted that in these cases, scientists decompose the mechanisms into both entities or parts and the activities or operations these perform, and also appeal to the ways in which these components are organized. Thus, Machamer et al. (2000) offer the following characterization of mechanisms: “*Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions.*” The characterization Bechtel and Abrahamsen (2005) advanced is quite similar: “A mechanism is a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism is responsible for one or more phenomena.” The differences in terminology (parts vs. entities; operations vs. activities) are not important for purposes here, but a significant point of disagreement concerns Machamer et al.’s contention that organization is sequential “from start or set-up to finish or termination conditions.” While they acknowledge that biological organization does sometimes involve forks and cycles, they nonetheless require start and termination conditions in their accounts and characterize the activities in between sequentially. Bechtel and Abrahamsen speak more generally of “orchestrated functioning” and have focused on examples, such as circadian rhythms, in which simultaneous operations in multiple feedback loops are crucial to the functioning of the mechanism. Noting the importance of computational modelling and the use of dynamical systems theory to understand how these mechanisms operate, they have proposed a framework of *dynamic mechanistic explanation* that incorporates dynamical analysis of the recomposed mechanism with the traditional focus on decomposition into parts and operations (Bechtel and Abrahamsen 2011; Bechtel 2011).

Mechanistic explanations seek to characterize the mechanism responsible for a given phenomenon. Sometimes such accounts of mechanisms are presented in linguistic descriptions. But often scientists find it helpful to represent the mechanism visually in diagrams that take advantage of two dimensions to situate parts and use conventions such as shape or color of icons to distinguish types of parts and

to indicate their operations. Arrows, with different styles of arrows representing different operations, portray the operations by which particular parts affect others. I will present some diagrams of mechanisms below, but for now I note how such use of diagrams distinguishes mechanistic explanations from nomological ones. Whereas nomological explanation insists on representing all information in propositions, researchers trying to understand mechanisms seek representational formats that support their reasoning about the mechanism. This links to another difference: nomological and mechanistic explanations provide alternative accounts as to how what is offered in explanation relates to the phenomenon being explained. In nomological explanations, logical derivations connect the laws to linguistic descriptions of the phenomena. Diagrams do not figure in logical derivations (although many logicians since Euler and Venn have appealed to diagrams to represent logical relationships), so the process of relating the explanations to the phenomena to be explained must be different on the mechanist account. Instead of drawing inferences, scientists simulate the operation of a mechanism to understand how it could generate the phenomenon. In some cases this can be done by mentally rehearsing each of the operations, often supported by a diagram or mentally imagining the mechanism as it would be presented in a diagram. As they confront more complex mechanisms with multiple simultaneous operations, scientists are increasingly appealing to animations to illustrate how a mechanism generates a phenomenon (McGill 2008).

Two weaknesses of both mental simulations and animations is that one may take the components of a mechanism as being able to do more than in fact they can or one may neglect important consequences of the operations of the components. This is particularly a risk when many operations are occurring simultaneously in the actual mechanism and these operations interact with each other in non-linear ways. Accordingly, researchers often appeal to model systems that they have constructed to emulate the way the parts operate or to computational models in which they characterize operations mathematically, to determine what the actual mechanism will do.

Philosophers of science adopting the nomological framework have traditionally eschewed trying to understand scientific discovery, insisting that the tools of logic could illuminate the evidential support for laws but not the processes by which they were discovered (Reichenbach 1938). (Some cognitive scientists, though, have ventured where philosophers feared to tread and advanced accounts of how such laws could be discovered. See Holland et al. 1986; Langley et al. 1987). When the project of explanation is understood as the discovery of mechanisms, philosophers are in a position to articulate the strategies through which scientists make discoveries. One approach is to focus on the reasoning strategies scientists use as they try to piece information together to develop an account of a mechanism (Bechtel and Richardson 1993/2010; Darden 2006; Darden and Craver 2002). Another is to focus in detail on how, in their experimental work, scientists intervene and manipulate biological mechanisms to elicit information about their parts and operations (Bechtel 2006; Craver 2002, 2007). Less has been done to date on the strategies through which scientists recompose mechanisms, especially in computational models, and use the results either to guide further experiments or revisions in the proposed mechanism (Bechtel and Abrahamsen 2010), but this is a topic ripe for additional philosophical research.

3 Delineating Phenomena

Mechanisms are invoked to explain phenomena and so it is important to specify what phenomena are. While many accounts of explanation have assumed that scientists try to explain observations (data), Bogen and Woodward (1988) convincingly demonstrate that what they in fact seek to explain are phenomena in the world. Although some phenomena only occur once, most are repeatable occurrences for which one can seek to specify the conditions under which they occur (doing so often requires considerable experimental inquiry). Examples of biological phenomena include the growth of plants, the births of mammals, circulation of blood, the metabolism of sugars, fats, and proteins to procure energy, and the conduction of an electrical signal down a nerve. These are what scientists seek to explain, not the data that provide evidence for them. Only when an observation or an experiment produces what is regarded as anomalous data do scientists turn their attention to explaining the data themselves.

An important part of delineating phenomena is to develop appropriate representations of them. Many of the philosophers advancing mechanistic explanation have focused on linguistic descriptions of phenomena (e.g., “proteins are synthesized by constructing strings of amino acids in the order specified in a sequence of DNA”). However, scientists are often interested in explaining much more specific features of phenomena, such as the rates at which a process occurs, and so characterize phenomena in terms of numerical values determined by empirical research. Frequently the numerical data they collect to characterize a phenomenon is presented in tables. However, because what they are generally interested in is the pattern exhibited in the numerical values, researchers must employ other representational techniques that make the pattern apparent. One way of doing this is in terms of equations, such as Weber’s (1834) psychophysical law identifying a constant proportion between a just noticeable change in a stimulus (ΔI) and the total quantity of the stimulus(I):

$$\frac{\Delta I}{I} = k.$$

Although this relation is often called a law, it does not play the role in explanation identified in nomological accounts: it specifies a relation for which an explanation is sought but does not explain its instances. It is important to distinguish between laws that offer explanations (Newton’s force laws were intended to explain the motion of bodies) and those such as Weber’s law that characterize phenomena. The latter still require explanation, either in terms of other laws or in terms of mechanisms that explain why the regularities hold. Weber himself proposed a possible mechanism, but a satisfactory explanation has not yet been advanced. Moreover, in the meantime researchers have advanced alternative mathematical representations, such as Stevens’ power law, that they claim to better characterize the relation between physical stimuli and how they are perceived. Often it proves difficult to develop an equation to characterize the phenomenon, and researchers instead


develop diagrammatic formats (such as illustrated below) to illustrate the pattern for which an explanation is sought.

Philosophers have typically presented the phenomenon to be explained as identified by observations. Even when phenomena are identified through observations (e.g., astronomical observations that were used by Copernicus, Kepler, Galileo, and Newton to characterize the behavior of planets), complex instruments and procedures are frequently required to secure the data from which the phenomenon can be elicited. In addition, in fields such as biology, scientists must intervene in nature to elicit the pattern for which explanation is developed. Even observational techniques, such as microscopic observation, require intervention to prepare specimens for observation through a microscope: water is removed from the specimen and it is chemically modified by stains and fixatives. What one observes in the microscope is the product of these manipulations, which are often quite brutal. An important question that biologists must address is whether they have generated an artifact or have accurately portrayed the phenomenon.²

I will illustrate the process of delineating the phenomenon in the case of circadian rhythms. The daily oscillations in some activities of living organisms are easily detectable if someone takes the effort to carry out observations across different periods of a day (the sleep-wake cycle in animals, phosphorescence in *Gonyaulax*, and the folding and unfolding of leaves in some plants). But others require instruments to detect them (e.g., the cycle in body temperature, of cell division in animals, or of gene expression). One of the key features of the phenomenon of circadian rhythms is that they are endogenously generated—they are not simply responses to the oscillations in daylight, temperature, etc., in the environment. To demonstrate that an oscillation in some feature or behavior of an organism meets this condition, researchers must set up conditions in which no likely environmental cue (known as a *Zeitgeber*) is available to set the phase of the oscillation. Thus, when de Mairan (1729) suspected the opening and closing of leaves in the *Mimosa* plant that he had observed was controlled endogenously, he had to resort to raising plants in continuous darkness and show that they continued to fold and unfold their leaves. However, it was always possible that the organism was responding to some other *Zeitgeber* that was itself responsive to the time of day. A crucial source of evidence in showing that rhythms were endogenously controlled was that when *Zeitgeber*s were removed, they oscillated with a period slightly different from 24 h, which should not be the case if they were responses to cues that were responsive to the day-night cycle of

²I have explored the challenge in determining whether evidence is an artifact or reflective of the real phenomenon and the strategies biologists use to address this challenge in the context of modern cell biology in Bechtel (2006). Since typically there is much that is unknown about how the procedures used actually work, researchers rely on such considerations as whether the evidence exhibits a distinct pattern distinguishable from noise, whether it can be at least partially corroborated using other techniques, and whether there are compelling theoretical explanations of the putative phenomenon. Although we commonly think of evidence as more secure and foundational to the explanatory hypotheses advanced, in fact evidence is often just as contested in science and evaluated in conjunction with the explanations offered.

Fig. 1 A basic actogram in which the *top bar* indicates a normal *light–dark* condition for the first 7 days and *constant darkness* for subsequent days. The *grey arrow* identifies the day a light pulse was administered. (From <http://www.photosensorybiology.org/id16.html>)



our planet (The name *circadian* reflects this: *circa* [about]+*dies* [day]). This required finding and representing patterns of behavior that oscillated with a period somewhat different from 24 h. The actogram format was developed to make this phenomenon manifest in a diagram.

An actogram is typically constructed from an automatic recording of an animal's activity. In perhaps the first actogram, Johnson (1926) devised a rotating disc on which a deflection was recorded every time a mouse moved. More recently, rodent activity is recorded from every rotation of a running wheel provided to the animal. As illustrated in Fig. 1, the time of day is represented across the top and successive days are represented by successive rows. A hash mark or other indicator represents the time when the animal rotated the wheel. Conventionally, the horizontal bars at the top show the periods during which the animal is exposed to light (the top bar shows that the animal was exposed to light from 4 to 16 h during the first 7 days, and subsequently kept in constant darkness). When information is represented in this format and one has learned the conventions that are employed, one can easily distinguish patterns in the animal's behavior. In this case, the animal began its activity somewhat earlier each day, revealing that its endogenous period is somewhat less than 24 h. This period when no Zeitgebers are present is known as free-running. Actograms can also be used to show how various perturbations affect free-running behavior. In this case, a pulse of light was presented 4 h after activity onset on the day indicated by the grey arrow. The actogram shows how this reset the animal's circadian rhythm, inserting a phase delay into what was otherwise a continuing pattern of phase advance due to constant darkness.

Varying the periods of light and darkness to which an animal is exposed is an effective way to explore the features of circadian phenomena. As already suggested in the previous actogram, when animals are exposed to relatively normal light–dark cycles, or to total darkness, they exhibit circadian rhythms. Individual pulses of

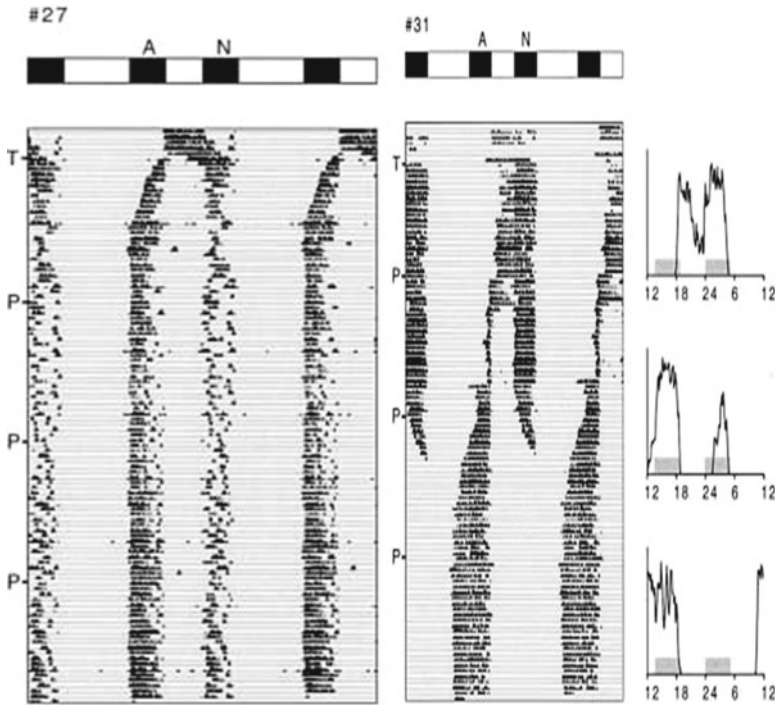


Fig. 2 Two actograms from Gorman (2001) showing the effects on the locomotor behavior of hamsters as a result of exposure to an unusual *light–dark* cycle

light while maintained under constant darkness can reset the circadian clock. But how do animals respond to atypical light–dark cycles? If the light–dark period is either very short (e.g., 19 h) or very long (e.g., 29 h), they typically become arrhythmic, which is manifest in an actogram by scattered bouts of activity on each day. More exotic are conditions under which an animal is exposed to two light and two dark periods each day. Gorman (2001) explored an arrangement in which hamsters were exposed to an arrangement of 9 h of light, 5 h of darkness, 5 h of light, and 5 h of darkness. In the actograms shown in Fig. 2, under these conditions hamsters would often develop split rhythms so that they were active during both dark phases. Sometimes, as illustrated in the actogram on the left, this would occur shortly after being introduced to the unusual lighting conditions (indicated by T on the left margin); in other cases, as illustrated in the actogram on the right, the splitting was delayed and not sustained.

An actogram is a particular diagrammatic format that circadian rhythm researchers have developed to represent a phenomenon of interest to them—changing patterns of behavior under different light–dark conditions. They are not, though, the only format employed—to portray how organisms are affected by light at different phases of their circadian period, researchers have developed phase-response curves that exhibit how much a particular form of light exposure advances or delays the phase

(C. H. Johnson 1999). Different fields of biology focus on different phenomena, and different diagrammatic formats have been developed and found effective for representing the phenomena of interest to them. Two general points should be noted. First, it is biologists themselves who must develop appropriate representational devices. Typically, in a given field the format is developed over time as researchers revise initial attempts until they arrive at a perspicuous format. Second, those using the diagrams must learn the conventions and this often takes time. After one has learned the technique, such diagrams seem transparent—they directly reveal the phenomenon. But students, for example, who have not yet become adept with the technique, fail to understand the phenomenon the diagram is illustrating.

4 Identifying and Decomposing a Mechanism

Once a phenomenon has been characterized,³ the challenge for researchers is to identify the mechanism and decompose it into its parts and operations. Although often researchers can identify the mechanism before identifying its parts and operations, sometimes the first clue to the mechanism results from identifying one of its parts. Circadian rhythm research illustrates both scenarios. In mammals researchers (Stephan and Zucker 1972; Moore and Eichler 1972; Moore and Lenn 1972) located the mechanism in the SCN both by showing that lesions to the SCN result in animals becoming arrhythmic and by tracing fiber tracts from the retina to the SCN (it was assumed that the clock must have input from the eyes if it were to be synchronized with the local light–dark cycle). In fruit flies, on the other hand, the first link to the mechanism was provided by identifying a gene, *period* (*per*), in which mutations result in flies with short or long endogenous rhythms or which are arrhythmic (Konopka and Benzer 1971). The brain locus of circadian control in flies (a small population of lateral and dorsal neurons) was only identified on the basis of studying the expression patterns of *per* (Helfrich-Förster 1996).

In either identifying the mechanism itself, or one of its parts, the strategies researchers used were much the same—they were seeking some entity whose operation correlates with the phenomenon of interest or that connects with entities to which the phenomenon is related. Perhaps the most common means of showing such a correlation is to ablate or mutate an entity and show a corresponding deficit in the phenomenon of interest. Studies ablating the SCN in mice and demonstrating loss of circadian rhythmicity exemplify this approach as do the mutation studies in

³Initial characterizations are often revised in the course of developing an explanation for them as that inquiry may reveal aspects of the phenomena that were not appreciated at the outset. In Bechtel and Richardson (1993/2010) we speak of this as reconstituting the phenomenon. Sometimes the reconstitution is quite major. For example, after 100 years of attempting to explain the phenomenon of animal heat, where such heat was viewed as an energy source that could support animal activity, it was recognized that such heat is actually a waste product and that the phenomenon of interest really involved the synthesis of adenosine triphosphate (ATP).

fruit flies. The case was made even more compelling when subsequent researchers showed that many circadian functions could be recovered when the SCN from another animal was placed in the third ventricle of the brain of the lesioned animal (Ralph et al. 1990). Other ways of generating such a correlation are to stimulate an entity and show an increase in the phenomenon of interest and to measure activity in the entity while the phenomenon of interest is occurring.⁴ While these various techniques show that the entity that is manipulated or recorded from is related to the phenomenon, they leave open the question whether on the one hand it is actually the responsible mechanism or a part within it or on the other hand only involved in a related activity. Removing the carburetor from a car or altering its operation affects whether and how the car will generate locomotion, but does not show it to be the mechanism responsible for the conversion of chemical energy to mechanical energy that explains locomotion. Typically, this latter question is only addressed by further developing the account of how the mechanism operates.

If research has successfully identified the mechanism itself, then subsequent research is directed at decomposing it into its parts and operations. In the case of the SCN, the immediate parts are the individual neurons and supporting glia cells. Initially, however, these were passed over as researchers proceeded directly to decomposing individual cells to identify the relevant constituents within them. As this research quickly dovetailed with research on fruit flies that had begun by identifying component parts of a mechanism, I turn first to how decompositional research proceeds once a part has been identified. As the idea of decomposing a mechanism into multiple parts suggests, a given part can only produce the phenomenon if it interacts with other parts. The quest is to identify these other parts and what they do. One approach is to iterate the first strategy, finding other components whose manipulation affects the phenomenon using the same strategies noted above, and once one or more additional components is identified, investigate what operation each performs. Another is to figure out what operation the part identified first performs and advance hypotheses about the other operations that are required to generate the phenomenon.

Research on circadian rhythms illustrates both strategies, but began with considering what operations can be attributed to *per*. Genes have effects when transcribed into mRNA and translated into proteins, and had researchers known what protein was expressed from *per* they might have begun by considering the reactions in which it could participate (proteins are frequently enzymes that catalyze reactions, and one might have considered whether the protein PER catalyzed reactions that relate to circadian function). Initially no known proteins were linked to *per* and it was only with the advent of cloning technology that researchers could begin to characterize the protein in terms of amino acid sequences that partially constituted it. This, however, provided a clue that focused researchers on a different way

⁴As with securing evidence about the phenomenon itself, these techniques involve manipulations, sometimes severe, and raise questions whether the information that is procured is just an artifact of the experimental manipulation.




Fig. 3 A representation of the mechanism proposed by Hardin et al. (1990) to explain circadian oscillations in fruit flies

to understand *per*'s contribution. Using cloning techniques, Hardin et al. (1990) determined that both *per* mRNA and PER oscillated with a circadian period, with the mRNA peaking several hours before that of the protein. This led them to suggest a mechanism consisting of three central parts: *per*, *per* mRNA, and PER. On their proposal, each performed a different operation: When uninhibited, *per* was transcribed into its mRNA. The mRNA in turn was transported from the nucleus to the cytoplasm, where it was translated into PER. Finally, PER was continually subject to degradation, but when enough accumulated, it was transported back into the nucleus where it inhibited *per* transcription (Fig. 3). At this point no new PER was produced, and as the existing molecules degraded, inhibition ceased and transcription began again.

While Hardin et al. proposed a mechanism schema that seemed to possess productive continuity, there were in fact numerous gaps. One concerned the manner in which PER degraded. The timing of that operation is crucial to generating an approximately 24 h oscillation. Initially of even greater concern was the question of how PER inhibited transcription of *per*. As the structure of PER was revealed, it seemed to lack a DNA binding region, suggesting the existence of another part which mediated between PER and *per*. Such a part was found when research on the mammalian clock mechanism revealed a gene dubbed *Clock* whose mutants exhibited disrupted circadian rhythms (Vitaterna et al. 1994). The CLOCK protein does possess a DNA binding region and mammalian CLOCK was demonstrated to bind to the promoter region of fruit fly *per*. PER was then viewed as interacting with CLOCK so that it no longer activated the transcription of *per*. Soon thereafter researchers identified both a fruit fly homolog of mammalian *Clock* and three mammalian homologs of *per*, resulting in highly intertwined research on the clock mechanisms in flies and mice (Bechtel 2009). Research into new clock components exploded. Both PER and CLOCK were found to operate as parts of compound molecules known as *heterodimers* (PER with TIM in fruit flies or CRY in mammals and CLOCK with CYCLE in fruit flies and BMAL1 in mammals). Concentrations of CLOCK in fruit flies and BMAL1 in mammals were themselves shown to oscillate,

leading to investigations that identified transcription factors that regulated their expression. As well, a number of kinases were identified that figure in the degradation of PER and TIM or CRY. As these parts were identified, researchers also established the immediate operations in which each was involved so that there is now a large catalog of known parts and operations of the circadian clocks of fruit flies and mammals.

Mechanistically oriented biologists have developed a host of tools for identifying component parts of mechanisms and determining what they do. Traditional tools, such as inducing mutations and registering concentrations of mRNAs with Northern blots and proteins with Western blots were extremely time consuming, but newer techniques have greatly simplified the study of individual genes and proteins. Researchers can directly target specific genes to knock-out or to suppress their transcription with tools such as small interfering RNA. And the ability to knock-in genes such as luciferase, an oxidative enzyme that generates light in organisms such as fireflies, has enabled researchers to observe oscillation in tissue-culture preparations through changing luminescence. Moreover, whenever a putative new clock gene is discovered, one can compare its DNA sequence to sequences stored in gene databanks to discover homologues in other organisms whose operations may be easier to assess. The research has provided an enormous wealth of information about parts and operations. Before leaving the topic of decomposition, I should note that the process of decomposition can be iterated—just as researchers decomposed the circadian mechanism into genes and proteins, they could decompose them into nucleotides and amino acids, and decompose them in turn into their component atoms, etc. Some accounts of reduction emphasize the iteration of decomposition down to the most basic entities science has identified at a given time (Bickle 2003). But from the point of view of mechanistic explanation, that is not the goal.

Researchers decompose a mechanism into the parts of the mechanism that explain its behavior. Some researchers might in turn want to explain how the parts work, and then they need to decompose those parts. Recently, for example, researchers have begun to identify how PER inhibits its own transcription, revealing the presence of PSF (polypyrimidine tract-binding protein-associated splicing factor) as a component of the PER complex and determining that it recruits SIN3A to scaffold assembly of a transcription inhibition complex that deacetylates histones in the chromatin of the *per* gene, preventing transcription (Duong et al. 2011). But it is important to note that this mechanism explains a different phenomenon—the inhibition of *per*—not the original phenomenon of maintaining circadian rhythms (in which *per*'s inhibition of PER was a basic operation).

5 Recomposing and Situating the Mechanism

Acquiring the catalog of parts and operations is an important step in developing mechanistic explanations, but until investigators determine how the operations of parts affect other parts (those they operate on) in a coordinated fashion to produce


the phenomenon, they have not yet explained the phenomenon. A researcher is no more satisfied with the catalog of parts and operations than you would be if you ordered a new car and it arrived as a collection of parts without even directions for putting them together. Before you have a functioning car, you need to figure out how the parts should be organized and work together to produce a vehicle that one can drive. Determining how the parts are organized and their operations orchestrated in the generation of the phenomenon is what I refer to as *recomposition*. In the course of science, scientists don't wait until they have a complete catalog of parts and operations to try to recompose the mechanism; rather, as they discover parts and operations they try to figure out how they go together to produce the phenomenon. Often the efforts at recomposition reveal the existence of other parts and the operations they perform. Thus, Fig. 3 above already reflects an attempt to recompose the mechanism, and this effort made clear to researchers that PER could not directly inhibit *per* transcription and other parts and operations remained to be discovered.

Although one can recompose a mechanism by narrating the envisioned operation of its parts, as I did above for Hardin et al.'s hypothesized mechanism, scientists commonly find it valuable to recompose the mechanism in a diagram. A diagram not only serves to present one's conception of the mechanism, but also supports reasoning about the mechanism. Just as in diagrams of phenomena, researchers can often see patterns in diagrams of mechanisms that would not otherwise be apparent.⁵ But there is yet another advantage of working with diagrams—a researcher can manipulate the diagram in the search for an alternative arrangement of operations that may better account for the phenomenon. This sometimes involves identifying gaps in the proposed account of the mechanism that need to be filled in. In this regard, it is interesting to consider how Hardin et al. themselves diagrammed the mechanism they were proposing (Fig. 4). While they clearly presented the idea of a feedback loop, they showed three alternative pathways by which the inhibition might arise (from the protein itself, from a product of a further reaction involving the protein, or from a behavior of the organism resulting from the protein) and two points at which it might effect the process of gene expression (transcription or translation). They also inserted question marks to indicate places where yet additional parts or operations might figure. Such a diagram becomes part of the reasoning processes of scientists as they considered both whether the proposed mechanisms could produce the phenomenon and how one might gain evidence for or against various hypotheses.

A diagram is a static representation that does not reveal how the operations are actually coordinated in the production of the phenomenon. With relatively simple mechanisms, especially ones that operate sequentially, mentally rehearsing the operations, perhaps guided by a narrative text, suffices to show how the mechanism

⁵For example, by comparing diagrams that have been developed for the circadian mechanisms identified in cyanobacteria, fungi, plants, and animals, researchers can readily see that although many of the parts are different, the overall organization is remarkably similar. This is in turn inspiring further inquiry that is revealing that even when the clocks contain different proteins, the domains and motifs that are crucial for the operations they perform are remarkably similar and may well have been evolutionary conserved (Stuart Brody, personal communication, January 2012.)

Fig. 4 Hardin et al.'s (1990) diagram of the mechanism they proposed to explain circadian rhythms in fruit flies in which they included *question marks* to indicate alternative places from which feedback might originate and at which it might terminate



will work. If the phenomenon is characterized quantitatively, though, it may be necessary to characterize each operation in equations and show mathematically that the phenomenon results. This becomes even more necessary when the mechanism involves non-sequential organization and the operations are non-linear. Here the ability to imagine the operation of the mechanism becomes highly unreliable. Humans are prone to fail to follow one of multiple effects of an operation and are poor at anticipating the results of non-linear interactions.

The circadian example makes this clear. Negative feedback was known to be a mechanism that could generate oscillatory behavior, but not all negative feedback processes generate sustained oscillations: If all the operations are linear, a feedback mechanism will settle into an equilibrium state in which each operation is equilibrated to the others—in the example I have been considering, just enough PER might be synthesized to maintain a constant level of suppression of synthesis to PER. A negative feedback system will only produce sustained endogenous oscillations if there are sufficient non-linear operations in the mechanism. Accordingly, one cannot determine by mental simulation whether a given feedback mechanism will generate sustained oscillations. The only options are either to construct such a mechanism (Elowitz and Leibler 2000) or to represent the operations of the mechanism in




Fig. 5 Equation (1) in Goldbeter’s (1995) model shown in relation to the relevant portion of Hardin et al.’s (1990) proposed circadian mechanism

mathematical equations and, through analysis or simulation, determine how the characterized mechanism will behave.⁶

Recognizing the problem of determining whether the mechanism proposed by Hardin et al. would sustain oscillations, Goldbeter (1995) represented an elaborated version of it that included operations that synthesized and degraded each component in a set of five differential equations. Each described the rate of change in the concentration of a substance in the mechanism (*per* mRNA, nuclear PER, etc.). Figure 5 shows how one of these equations describes operations that affect the concentrations of *per* mRNA—the transcription of *per* into *per* mRNA and the hypothesized degradation of *per* mRNA. The equation includes both variables (M for the concentration of *per* mRNA; P for the concentration of PER in the nucleus) and parameters (K_I and K_m represent the Michaelis constant and v_s and v_m the maximum rates for the two reactions respectively). Modelers choose values for the parameters both with an eye to making the simulation achieve the desired end and to characterizing the actual biological operations—they may speak of the parameters they employ as “biologically plausible.”

⁶Neither of these strategies is perfect. Mathematical modeling depends on developing equations that accurately describe the operations in the mechanism. Researchers avoid this problem in synthetic models, letting actual physical components operate as they do. Such models often produce quite surprising results, which then motivate the creations of new computational models in the attempt to explain them. See Danino et al. (2010) who provide a particular compelling example of the productive interaction between synthetic models and computational models in understanding synchronization of oscillations.

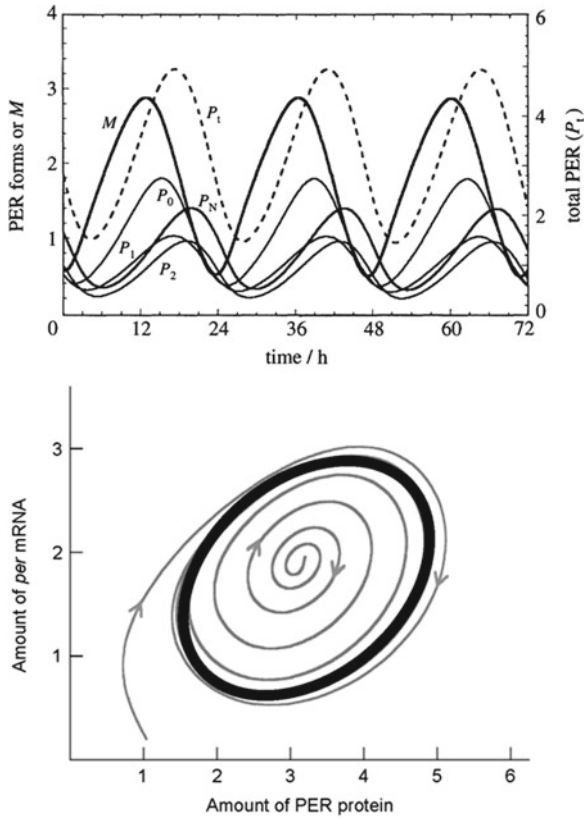


Fig. 6 Two diagrams Goldbeter (1995) used to illustrate the behavior of his computational model of Hardin et al.'s proposed mechanism. In the top diagram the changes in variables in his model (after initial transients) are plotted against time, whereas in the one on the *bottom* successive values (indicated by arrows) of two variables are plotted in phase space. This shows the transients as the mechanism approaches the limit cycle (*dark oval*), at which point it oscillates indefinitely

Sometimes the behavior of a mechanism can be deduced directly from the equations and specified parameter values, but typically this is not possible and modelers rely on simulations run on computers. A given simulation begins by assigning initial values to the variables and then iteratively applying the equations to determine values at subsequent times. Typically multiple simulations will be conducted to ascertain the behavior of the mechanism. The results of the simulation are a set of numbers, which modelers then plot in a diagram to decipher whether a pattern results. Goldbeter chose to plot the results in two ways. As shown at the top in Fig. 6, he first plotted time on the abscissa and values of variables on the ordinate, revealing how the values oscillated across time (once the simulation reached a stable oscillation). Often oscillatory patterns, especially before a stable oscillation is achieved, are more perspicuously shown in phase space, as illustrated at the bottom

of Fig. 6. Here the values of two variables, the total concentration of the PER protein and *per* mRNA, are indicated on the abscissa and ordinate and time is conveyed only in the succession of locations plotted (the order of successive points is indicated with arrows). In this case the phase space plot makes clear that from different initial values, the behavior will follow a trajectory towards the dark oval, referred to as a *limit cycle* (in this figure two trajectories are shown). If the values of PER and *per* mRNA are on the oval, they will continue to change so as to follow the oval, thus showing that the oscillations will be sustained indefinitely. This analysis illustrates the integration of mechanism and dynamics in dynamic mechanistic explanations.

I noted above that in the years after Goldbeter put forward his model, researchers discovered numerous other components of the intracellular clock mechanism. Since additional components could significantly alter the behavior of the mechanism, disrupting its ability to sustain oscillations, it was necessary to build more complex models to ascertain their behavior. Goldbeter together with Leloup developed models incorporating all the known components of the fruit fly (Leloup and Goldbeter 1998) and the mammalian (Leloup and Goldbeter 2004) clock mechanism. Other modelers have pursued different strategies, developing less detailed models that enable running simulated experiments that, for example, might reveal which components of the mechanism are most crucial for producing the phenomenon (Smolen et al. 2001; Relógio et al. 2011).

So far I have focused on recomposing the mechanism in diagrams and computational models, activities that are important to investigators' attempts to understand how mechanisms will actually function. In these efforts, however, researchers commonly abstract from the larger contexts in which the mechanism actually functions, but these can have important consequences for the functioning of the mechanism. Appreciating them requires not just recomposing the mechanism, but situating it in contexts in which it usually operates. Just as researchers often must iterate the process of decomposition, they often must also iterate the processes of recomposing and situating the mechanism since its operation may be affected not just by other activities in its local environment, but also by activities in the environment of the system of which it is a component. I will briefly indicate three such levels of situating and recomposing required to understand circadian rhythms.

Earlier I noted that in mammals researchers first located the circadian clock in the SCN, a relatively small region of the hypothalamus consisting, in mice, of 16,000–20,000 neurons on each side of the brain. As they pursued the question of how individual neurons in the SCN maintained time, researchers implicitly assumed each SCN neuron operated in the same manner. Welsh et al. (1995), however, discovered that when dispersed in culture both the phase and the period of oscillations varied substantially across neurons. More recent research has shown that many SCN neurons do not maintain rhythmicity when deprived of inputs from their neighbors, and some shift between normal and super-long periods (Meeker et al. 2011). These discoveries make it clear that the oscillations within individual neurons are modulated by the behavior of their neighbors. To understand this behavior researchers are turning more and more to computational modeling as providing the most tractable

way of investigating how such a complicated system might operate. These modeling efforts typically begin by using models constructed for the intracellular oscillators and adding terms and equations to them to represent hypothesized interactions between SCN neurons. For example, Vasalou et al. (2009) showed that by assuming the SCN had the structure of a small-world (a network structure in which most connections are local but a few extend long distances) they could capture many details of the phenomenon of SCN behavior. The assumption that the SCN is a small world is highly plausible given recent discoveries of the ubiquity of such organization in biological systems, but before one can assume that the model reveals the behavior of the actual mechanism, further research is required to determine how closely the actual organization of the SCN corresponds to what is proposed in the models.

The fact that lesioning the SCN in mice eliminated circadian behavior and that implanting SCN tissue from a mouse with a clock mutation into the third ventricle of a lesioned mouse would restore some circadian behaviors with the period found in the mutant convinced researchers that the SCN was the master clock (Ralph et al. 1990). One of the consequences of identifying the components of the central clock in the SCN was the discovery that the same genes are also expressed in most cells of the body and in them they also work together as clock mechanisms. Since timing in peripheral tissue was lost with lesions to the SCN, these peripheral clocks were thought to be slaves. This supported viewing the SCN as the largely autonomous locus of the clock mechanism, constrained only by inputs it received from Zeitgebers and sending outputs to other tissues. More recently it has been demonstrated that the clock mechanisms in peripheral tissues do not stop oscillating in the absence of the SCN but rather become desynchronized from each other. This has led to rethinking the relationship of the SCN to this other clocks: the SCN is better thought of as an orchestra conductor than a slave master (Davidson et al. 2004). Circadian oscillators in the liver have been shown to operate semi-independently of the SCN, being entrained to feeding schedules outside the organism's normal feeding times (Damiola et al. 2000). Moreover, there is increasing evidence that metabolic activities, many of which are regulated by the liver, have effects of the SCN (Nakahata et al. 2008). In addition, researchers are increasingly discovering avenues through which clocks in other organs of the body and the processes they regulate feed back on the SCN and affect its behavior, rendering it important to understand how the SCN is situated in the organism (Pezuk et al. 2010).

Finally, organisms with circadian rhythms operate in an environment that on this planet has a 24 h light–dark cycle. It has long been recognized that having an endogenous circadian clock is crucial for organisms to function in their environments as many activities must be performed at appropriate times of day and it often takes several hours for biological systems to prepare to carry out these activities (for example, enzymes need to be synthesized to perform photosynthesis in plants or to digest food in animals). For most organisms the light–dark cycle is externally provided and they must only entrain their endogenous clocks to it. Humans, however, have developed artificial environments in which their exposure to light and conduct of activities is dissociated from the light–dark cycle provided by the Earth's rotation. As a result, our endogenous circadian rhythms are confronted with

discordant entrainment signals from environments we have created. The experience of jet lag after travel across multiple time zones makes apparent the disruptions abruptly altered light–dark cycles can have until, over several days, our endogenous clock is re-entrained to the local light–dark cycle. The severe health effects of shift work, which results in our endogenous clock being confronted with a very unnatural light–dark cycle, are increasingly being identified (Maywood et al. 2006; Wang et al. 2011).

Recomposing and situating mechanisms is crucial for understanding their behavior. The result is often a much more holistic understanding of the phenomenon one seeks to explain. It should be noted, though, that appreciation of the whole is often dependent on first developing at least a basic understanding of the composition of the mechanism. Models of mechanisms generated before research reveals actual parts and operations are at best only hypotheses about how the mechanism might operate. Once accounts can be grounded on some understanding of parts and operations, then the task of understanding how the whole mechanism is actually organized and its functioning orchestrated is much more tractable (although certainly not easy). Moreover, once one has an account of how the recomposed mechanism might work, researchers can both evaluate the effects of different contexts on its operation and figure out how these effects are realized in the mechanism.

6 Conclusion

When biologists seek to explain a phenomenon, in many contexts what they are looking for is an account of the mechanism responsible for it. I have described some of the distinctive features of mechanistic research, including its frequent reliance on diagrammatic representations and the strategies for discovering mechanisms. A critical first step is to delineate the phenomenon for which explanation is sought. Although sometimes characterized in verbal descriptions (e.g., maintain circadian oscillations), typically phenomena are characterized quantitatively with diagrams used to represent the pattern elicited from quantitative information. The second step is to identify the responsible mechanism and decompose it into its parts and operations. This process of taking systems part—finding the part that constitutes the mechanism and discovering its components—is what distinguishes mechanistic research. I described how such research led to the discovery of the SCN as the locus of the central clock in mammals and the genes and proteins whose operations figure in maintaining oscillations.

As critical as decomposition is, however, it is equally important to recombine the mechanism so as to understand how the parts operate together. Scientists often present their understanding of the mechanism in diagrams, but to determine what will actually result from the parts performing their operations they turn to simulations. In simple cases, researchers can mentally simulate the mechanism of interest, but increasingly, as research reveals non-sequential organization involving nonlinear operations, biologists appeal to computational models. The result is what I term

dynamic mechanistic explanation. I illustrated how a relatively simple model was used to demonstrate that an early hypothesis about how parts interacted might in fact produce circadian oscillations. Beyond recomposing individual mechanisms, researchers often come to recognize ways in which the mechanism is affected by the context in which it operates, and this requires that they situate the mechanism in various environments and assess, often through computational models, how they will affect the mechanism's behavior.

Thinking in terms of mechanisms is quite intuitive, especially in cultures exposed to modestly complex technology. Most of us are familiar with taking mechanisms apart, either to diagnose problems and repair them or just out of curiosity as to how they work. But, as van Mil et al. (2013) discuss, students do not readily extend this perspective to biological mechanisms. Their diagnosis is in part that biology education focuses on the functions of mechanisms, not on how they work. Another part of the explanation is perhaps that our prototypes of machines are devices made out of solid materials (e.g., wood or metal) with relatively clearly delineated parts organized in a fairly straightforward manner. As these characteristics of prototypical machines are violated, as they often are in biology, people are less inclined to adopt a mechanistic perspective.⁷ Professional biologists have a several hundred-year history of adopting and enriching the mechanistic perspective (although one finds biologists repudiating the machine metaphor as they focus on the complex organization found in many biological systems). Moreover, they have become accustomed to thinking of components just in terms of the contributions they make to the whole mechanism and then seeking evidence as to their physical constitution.⁸ This strategy of starting with a hypothesized functional decomposition requires imagination (and reasoning by analogy to other known mechanisms) and may require cultivation.

While biologists have become enculturated into this extension of mechanical thinking to biology, it may not come naturally to students and the process of developing mechanistic explanations in biology may need to be explicitly articulated. The example of circadian rhythms, as developed in this chapter, drawing as it does on the analogy with manufactured clocks, can provide a helpful entrée for getting students to think mechanistically about biological processes. I have noted the value of diagrams in conveying an understanding of mechanisms. But diagrams themselves may require commentary that explicitly notes the parts and operations shown and helps students learn to rehearse these operations so as to understand how they work

⁷In fact, as machines become more complex and rely on electronic circuitry that is not readily decomposed, we relatively quickly abandon our mechanistic perspective on how they work and settle simply for learning how to control their operation.

⁸van Mil et al. (2013) identify the failure to understand how proteins change conformation as one factor in students' failure to think causally about protein interactions in the cell. While not denying the importance of conveying such an understanding of proteins, I suspect it is not the main problem in thinking causally about cell constituents as biologists themselves came to understand protein actions well before they understood their conformation. Rather, the problem seems more immediate—students are not encouraged to think about how component processes may work together in producing physiological effects.

together to produce the phenomenon. Spending the time needed to work through a diagram and to illustrate how one can use it to think through the workings of a mechanism may help bring students into the culture in which biologists operate.

As I have noted above, discovering or learning how a mechanism works is a reductionist inquiry—it requires decomposing a mechanism into its parts and the operations they perform. But, as I have emphasized, efforts at decomposition need to be complemented by research seeking to recompose a mechanism and to understand how it is situated. Understanding a mechanism requires dexterity in moving down and up between levels of organization. van Mil et al. (2013) note that thinking in terms of levels and moving between them is often difficult for students. Given the challenges that scientists have confronted in moving up and down levels in their inquiries, it is not surprising that students face a challenge. The problem is exacerbated by the fact that nature does not present itself well-delineated in terms of levels—it is only in the process of developing mechanistic explanations that researchers come to recognize both components and the often transient structures into which they are composed. Nonetheless, if students are to develop appropriate sophistication in understanding and developing mechanistic accounts, they need to develop facility in thinking about levels. Working through examples is often the best way to acquire such facility, and the circadian example developed here may prove useful in cultivating such facility.

Acknowledgments I thank Kostas Kampourakis and an anonymous reviewer for helpful comments and suggestions that improved this essay. Research for this chapter was supported both by a grant from the Dean of Arts and Sciences at UCSD and by the National Science Foundation (USA) under Project Number 1127640.

References

- Bechtel, W. 2006. *Discovering cell mechanisms: The creation of modern cell biology*. Cambridge: Cambridge University Press.
- Bechtel, W. 2009. Generalization and discovery by assuming conserved mechanisms: Cross species research on circadian oscillators. *Philosophy of Science* 76: 762–773.
- Bechtel, W. 2011. Mechanism and biological explanation. *Philosophy of Science* 78: 533–557.
- Bechtel, W., and A. Abrahamsen. 2005. Explanation: A mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 421–441.
- Bechtel, W., and A. Abrahamsen. 2010. Dynamic mechanistic explanation: Computational modeling of circadian rhythms as an exemplar for cognitive science. *Studies in History and Philosophy of Science Part A* 41: 321–333.
- Bechtel, W., and A. Abrahamsen. 2011. Complex biological mechanisms: Cyclic, oscillatory, and autonomous. In *Philosophy of complex systems*, Handbook of the philosophy of science, vol. 10, ed. C.A. Hooker, 257–285. New York: Elsevier.
- Bechtel, W., and R.C. Richardson. 1993/2010. *Discovering complexity: Decomposition and localization as strategies in scientific research*. Cambridge, MA: MIT Press. 1993 edition published by Princeton University Press.
- Bickle, J. 2003. *Philosophy and neuroscience: A ruthlessly reductive account*. Dordrecht: Kluwer.
- Bogen, J., and J. Woodward. 1988. Saving the phenomena. *Philosophical Review* 97: 303–352.

- Booger, F.C., F.J. Bruggeman, R.C. Richardson, A. Stephan, and H.V. Westerhoff. 2005. Emergence and its place in nature: A case study of biochemical networks. *Synthese* 145: 131–164.
- Craver, C.F. 2002. Interlevel experiments and multilevel mechanisms in the neuroscience of memory. *Philosophy of Science* 69: S83–S97.
- Craver, C.F. 2007. *Explaining the brain: What a science of the mind-brain could be*. New York: Oxford University Press.
- Damiola, F., N. Le Minh, N. Preitner, B. Kornmann, F. Fleury-Olela, and U. Schibler. 2000. Restricted feeding uncouples circadian oscillators in peripheral tissues from the central pacemaker in the suprachiasmatic nucleus. *Genes and Development* 14: 2950–2961.
- Danino, T., O. Mondragon-Palomino, L. Tsimring, and J. Hasty. 2010. A synchronized quorum of genetic clocks. *Nature* 463: 326–330.
- Darden, L. 2006. *Reasoning in biological discoveries: Essays on mechanisms, interfield relations, and anomaly resolution*. Cambridge: Cambridge University Press.
- Darden, L., and C. Craver. 2002. Strategies in the interfield discovery of the mechanism of protein synthesis. *Studies in History and Philosophy of Biological and Biomedical Sciences* 33: 1–28.
- Davidson, A.J., S. Yamazaki, and M. Menaker. 2004. SCN: Ringmaster of the circadian circus or conductor of the circadian orchestra? In *Molecular clocks and light signalling*, ed. D.J. Chadwick and J.A. Goode, 110–125. Chichester: Wiley.
- De Mairan, J.-J.d.O. 1729. Observation Botanique. *Histoire de l'Academie Royale Sciences*, pp. 35–36.
- Duong, H.A., M.S. Robles, D. Knutti, and C.J. Weitz. 2011. A molecular mechanism for circadian clock negative feedback. *Science* 332: 1436–1439.
- Elowitz, M.B., and S. Leibler. 2000. A synthetic oscillatory network of transcriptional regulators. *Nature* 403: 335–338.
- Goldbeter, A. 1995. A model for circadian oscillations in the *Drosophila* Period protein (PER). *Proceedings of the Royal Society of London B: Biological Sciences* 261: 319–324.
- Gorman, M.R. 2001. Exotic photoperiods induce and entrain split circadian activity rhythms in hamsters. *Journal of Comparative Physiology A: Sensory Neural and Behavioral Physiology* 187: 793–800.
- Griffiths, P.E., and R. Gray. 1994. Developmental systems and evolutionary explanation. *Journal of Philosophy* 91: 277–304.
- Hardin, P.E., J.C. Hall, and M. Rosbash. 1990. Feedback of the *Drosophila period* gene product on circadian cycling of its messenger RNA levels. *Nature* 343: 536–540.
- Helfrich-Förster, C. 1996. *Drosophila* rhythms: From brain to behavior. *Seminars in Cell and Developmental Biology* 7: 791–802.
- Hempel, C.G. 1965. Aspects of scientific explanation. In *Aspects of scientific explanation and other essays in the philosophy of science*, ed. C.G. Hempel, 331–496. New York: Macmillan.
- Holland, J.H., K.J. Holyoak, R.E. Nisbett, and P.R. Thagard. 1986. *Induction: Processes of inference, learning and discovery*. Cambridge, MA: MIT Press.
- Johnson, M.S. 1926. Activity and distribution of certain wild mice in relation to biotic communities. *Journal of Mammalogy* 7: 254–277.
- Johnson, C.H. 1999. Forty years of PRCs – What have we learned? *Chronobiology International* 16: 711–743.
- Konopka, R.J., and S. Benzer. 1971. Clock mutants of *Drosophila melanogaster*. *Proceedings of the National Academy of Sciences USA* 89: 2112–2116.
- Langley, P., H.A. Simon, G.L. Bradshaw, and J.M. Zytkow. 1987. *Scientific discovery: Computational explorations of the creative process*. Cambridge: MIT Press.
- Leloup, J.-C., and A. Goldbeter. 1998. A model for circadian rhythms in *Drosophila* incorporating the formation of a complex between the PER and TIM proteins. *Journal of Biological Rhythms* 13: 70–87.
- Leloup, J.-C., and A. Goldbeter. 2004. Modeling the mammalian circadian clock: Sensitivity analysis and multiplicity of oscillatory mechanisms. *Journal of Theoretical Biology* 230: 541–562.

- Machamer, P., L. Darden, and C.F. Craver. 2000. Thinking about mechanisms. *Philosophy of Science* 67: 1–25.
- Maywood, E.S., J. O'Neill, G.K.Y. Wong, A.B. Reddy, and M.H. Hastings. 2006. Circadian timing in health and disease. *Progress in Brain Research* 153: 253–269; Elsevier.
- McGill, G.I. 2008. Molecular movies ... Coming to a lecture near you. *Cell* 133: 1127–1132.
- Meeker, K., R. Harang, A.B. Webb, D.K. Welsh, F.J. Doyle, G. Bonnet, et al. 2011. Wavelet measurement suggests cause of period instability in mammalian circadian neurons. *Journal of Biological Rhythms* 26: 353–362.
- Moore, R.Y., and V.B. Eichler. 1972. Loss of a circadian adrenal corticosterone rhythm following suprachiasmatic lesions in the rat. *Brain Research* 42: 201–206.
- Moore, R.Y., and N.J. Lenn. 1972. A retinohypothalamic projection in the rat. *The Journal of Comparative Neurology* 146: 1–14.
- Nagel, E. 1961. *The structure of science*. New York: Harcourt, Brace.
- Nakahata, Y., M. Kaluzova, B. Grimaldi, S. Sahar, J. Hirayama, D. Chen, et al. 2008. The NAD⁺-dependent deacetylase SIRT1 modulates CLOCK-mediated chromatin remodeling and circadian control. *Cell* 134: 329–340.
- Nakajima, M., K. Imai, H. Ito, T. Nishiwaki, Y. Murayama, H. Iwasaki, et al. 2005. Reconstitution of circadian oscillation of cyanobacterial KaiC phosphorylation in vitro. *Science* 308: 414–415.
- Pezuk, P., J.A. Mohawk, T. Yoshikawa, M.T. Sellix, and M. Menaker. 2010. Circadian organization is governed by extra-SCN pacemakers. *Journal of Biological Rhythms* 25: 432–441.
- Pigliucci, M., and M. Boudry. 2011. Why machine-information metaphors are bad for science and science education. *Science & Education* 20: 453–471.
- Post, E.L. 1936. Finite combinatorial processes – Formulation I. *Journal of Symbolic Logic* 1: 103–105.
- Ralph, M.R., R.G. Foster, F.C. Davis, and M. Menaker. 1990. Transplanted suprachiasmatic nucleus determines circadian period. *Science* 247: 975–978.
- Reichenbach, H. 1938. *Experience and prediction*. Chicago: University of Chicago Press.
- Relógio, A., P.O. Westermark, T. Wallach, K. Schellenberg, A. Kramer, and H. Herzl. 2011. Tuning the mammalian circadian clock: Robust synergy of two loops. *PLoS Computational Biology* 7: e1002309.
- Smolen, P., D.A. Baxter, and J.H. Byrne. 2001. Modeling circadian oscillations with interlocking positive and negative feedback loops. *Journal of Neuroscience* 21: 6644–6656.
- Stephan, F.K., and I. Zucker. 1972. Circadian rhythms in drinking behavior and locomotor activity of rats are eliminated by hypothalamic lesions. *Proceedings of the National Academy of Sciences USA* 69: 1583–1586.
- Turing, A. 1936. On computable numbers, with an application to the Entscheidungs problem. *Proceedings of the London Mathematical Society Second Series* 42: 230–265.
- van Mil, M., D. Boerwinkel, and A. Waarlo. 2013. Modelling molecular mechanisms: A framework of scientific reasoning to construct molecular-level explanations for cellular behaviour. *Science & Education* 22(1): 93–118.
- Vasalou, C., E.D. Herzog, and M.A. Henson. 2009. Small-world network models of intercellular coupling predict enhanced synchronization in the suprachiasmatic nucleus. *Journal of Biological Rhythms* 24: 243–254.
- Vitaterna, M.H., D.P. King, A.-M. Chang, J.M. Kornhauser, P.L. Lowrey, J.D. McDonald, et al. 1994. Mutagenesis and mapping of a mouse gene, *Clock*, essential for circadian behavior. *Science* 264: 719–725.
- Wang, X.-S., M.E.G. Armstrong, B.J. Cairns, T.J. Key, and R.C. Travis. 2011. Shift work and chronic disease: The epidemiological evidence. *Occupational Medicine* 61: 78–89.
- Weber, E.H. 1834. *De pulsus, resorptione, auditu et tactu: Annotationes anatomicae et physiologicae*. Leipzig: Koehler.
- Welsh, D.K., D.E. Logothetis, M. Meister, and S.M. Reppert. 1995. Individual neurons dissociated from rat suprachiasmatic nucleus express independently phased circadian firing rhythms. *Neuron* 14: 697–706.