

Explaining Features of Fine-Grained Phenomena Using Abstract Analyses of Phenomena and Mechanisms: Two Examples from Chronobiology

William Bechtel

Department of Philosophy and Center for Circadian Biology
University of California, San Diego

Abstract

Explanations of biological phenomena such as cell division, protein synthesis or circadian rhythms commonly take the form of models of the responsible mechanisms. Recently philosophers of science have attempted to analyze this practice, presenting mechanisms as organized collections of parts performing operations that together produce the phenomenon. But in some cases what researchers seek to explain is not a general phenomenon, but a specific feature of a more fine-grained phenomenon. In some of these cases, it is not the model of the mechanism that performs the explanatory work. I consider a case in which the investigator offered an abstract representation of a fine-grained phenomenon to show why it had the feature in question. I consider a second case in which a researcher abstracted from the mechanism to identify a design principle that explains why the functioning mechanism exhibits a specific feature.

I. Introduction

A common philosophical account of explanation in biology is that biologists advance models of mechanisms to explain general biological phenomena such as cell division, protein synthesis, or circadian rhythms. On this account, mechanisms are construed as organized collections of parts performing operations that together produce the phenomenon (Bechtel & Richardson, 1993/2010; Bechtel & Abrahamsen, 2005; Machamer, Darden, & Craver, 2000; Craver & Darden, 2013). This picture fits textbook presentations of biology and many review papers, which seek to explain general phenomena. But often, especially in research studies, the focus is on a specific feature of a more fine-grained phenomena, not the general phenomenon itself. What scientists offer as explanations of these features of more narrowly construed phenomena also differs from the explanations for general phenomena. Unlike in the case of general phenomena, models of mechanisms often fail to provide the sought after explanation.

In what follows I will first introduce in section 2 the claim that often the target of explanations is a specific feature of a fine-grained phenomena. In subsequent sections I describe two examples, both drawn from research on the general phenomenon of circadian rhythms, in which researchers focused on particular features of fine-grained phenomena. Rather than offering a detailed mechanistic account of the responsible mechanism, the researchers adopted different strategies. In the first example, discussed in section 3, Winfree, working before even a modestly articulated mechanism for circadian rhythms had been proposed, appealed to a mathematical truth (Lange, 2012) both to predict an

important and surprising feature of the specific fine-grained phenomenon of entrainment of circadian rhythms. From the mathematical truth and the empirical details of entrainment, he offered an explanation of why the phenomenon had to have that feature. The second example, discussed in section 4, involves research after a relatively complicated mechanism had been proposed. While accepting that the proposed mechanism could generate circadian rhythms, many researchers were puzzled about why it produced sustained oscillations. To address this issue, Ueda abstracted from the details of the mechanism to characterize two design principles (Green, 2015)¹ whose ability to generate sustained oscillations was already well understood. In both of these cases, while an account of a mechanism remains relevant to understanding the general phenomenon, it is not what explains the features of the fine-grained phenomenon on which the researchers were focused.

To provide background for the specific cases that follow, I offer here a brief sketch of the history of research on circadian rhythms. Observations of daily rhythms in behavior or physiological activities date back to ancient Greece, with the daily folding and unfolding of plant leaves providing one of the most commonly cited examples. By placing mimosa plants in a dark closet and observing that they continued to fold and unfold their leaves, De Mairan (1729) provide some of the first evidence that these rhythms were generated endogenously. Suspicions remained that some external cue was responsible for such behavior until crucial evidence was provided that when kept isolated from known environmental time cues (*Zeitgebers*) rhythms continued, but with a period slightly different than 24 hours (a condition referred to as *free-running*). To emphasize that the period is only approximately 24 hours, Halberg (1959) introduced the name *circadian* (from *circa* and *dies*). Researchers at the time of the 1960 Symposium on Biological Clocks at Cold Spring Harbor, the first large conference of circadian researchers, largely accepted the hypothesis that these rhythms were endogenously generated. The use of the term *clock*² in the title alludes to a growing interest in identifying the mechanism responsible for these rhythms. Although many speculative proposal were advanced (Edmunds, 1988), progress in procuring relevant empirical evidence was slow. Through a forward screen of mutant fruit flies with altered rhythms, Konopka and Benzer (1971) identified a gene which, when mutated, resulted in slow or fast rhythms or arrhythmic behavior, which they named *period* or *per*. Studies of the operations in which *per* participated had to await the advent of cloning in the 1980s. Using cloning, Hardin, Hall, and Rosbash (1990) discovered that the mRNA and protein PER into which *per* was transcribed and translated both also oscillated with about a four-hour phase delay between them, and hypothesized a transcription-translation feedback loop (TTFL) in which the protein PER feeds back to inhibit the transcription of its own gene (Figure 1). Mental animation of such a feedback loop provides an intuitive account of how a TTFL could generate circadian oscillation:

¹ Levy and Bechtel (2013) explore how abstracting from details of a mechanism can reveal the organization of the mechanism that is responsible for certain phenomena. In this paper I extend the focus on abstraction as a tool for developing explanations further.

² Although the clock metaphor was actually introduced by Brown, who was one of the last holdouts for the view that circadian rhythms depended on environmental cues, Pittendrigh soon adopted it to characterize the endogenous mechanism he took to be responsible.

when PER concentrations are low, transcription will occur, resulting in a steady increase in the concentration of PER. As the concentration of PER rises, however, it exerts an inhibitory effect on transcription. Concentrations will then drop as PER gradually degrades and is not replaced. As the concentration of PER drops, though, PER ceases to inhibit transcription and concentrations of PER will raise again. Research over the course of the next two decades filled in this initial proposal resulting in the complicated mechanism shown in Figure 8 below.




Figure 1. The delayed negative feedback mechanism for generating circadian rhythms proposed by Hardin et al. (1990).

The articulation of this mechanism for the general phenomenon of circadian rhythms is rightly regarded as a major achievement. While it provides a backdrop, it does not itself explain the features of the more fine-grained circadian phenomena on which I will focus in sections 3 and 4.

2. Explaining Specific Features of Fine-Grained Phenomena

In presenting phenomena as the targets of explanation, mechanists have appealed to Bogen and Woodward's (1988) contention that scientific theories do not explain observations or data but phenomena. Phenomena of the sort that are explained in science are not individual occurrences from which data might be procured, but repeatable ones. Bogen and Woodward offer as examples of phenomena "weak neutral currents, the decay of the proton, and chunking and recency effects in human memory." When mechanists adopt Bogen and Woodward's distinction, they took up biological examples such as protein synthesis, the generation of action potentials, oxidative phosphorylation, and circadian rhythms. As in Bogen and Woodward's examples, researchers had to carry out a great deal of research and then extract from the conflicting and noisy data they generated a proposal as to the general phenomenon that was occurring. Statistical and data reduction techniques figured significantly in this process along with designing experiments to rule out possible confounds. This research was not itself directed at explanation, but at accurately characterizing phenomena.

Often in the early stages of inquiry, phenomena are characterized very generally. Phenomena such as protein synthesis and circadian rhythmicity occur widely amongst living organisms and the way they are characterized abstracts from particular details so as to capture important common features. For example, circadian rhythms are characterized

as endogenously generated oscillations of approximately 24-hours that are manifest in many physiological and behavior activities, are entrainable in response to different environmental stimuli, and are temperature compensated. These features characterize circadian rhythms in organisms from cyanobacteria to humans even though there are many important differences that can be identified. Once characterized, many researchers pursue the project of trying to explain them by identifying and describing the responsible mechanism. For each of the biological phenomena noted above there are now generally accepted mechanistic accounts that describe the types of entities involved, the activities or operations they perform, and how the components are organized in space and time so as to produce the phenomenon. Accounts of these phenomena are presented in textbooks and sometimes in review articles.

These general phenomena are not the explanatory targets of individual research projects. Rather, individual research projects are directed at specific features attributed to phenomena that themselves are described in a much more fine-grained manner. If one hopes to understand explanation as it is pursued in biology and related fields, one needs to focus on specific features of more fine-grained accounts of phenomena. To appreciate one reason researchers focus on features of more fine-grained phenomena, it will be useful to consider how data contributes to characterizing phenomena. The data researchers collect in the course of delineating phenomena are often detailed and quantitative. As a result, one can describe phenomena in a much more detailed way, quantifying the phenomenon, not just describing it qualitatively. For example, circadian rhythms not only involve approximately 24-hour oscillations in physiology and behavior, but each of these activities exhibits a different 24-hour pattern (e.g., peaking at different times of day). And, as we will see in the next section, circadian oscillations are not just entrainable to environmental stimuli, but show a particular quantitative pattern of advancement or delay in response to specific types and amounts of stimuli. These specific patterns are then the target of explanatory research.

To characterize these features, researchers often find it critical to identify the phenomena themselves in a more fine-grained way by, for example, focusing only on the phenomenon as exhibited in a given group of organisms. While recognizing that there will be differences between model organisms (such as mice) and the organisms of interest (e.g., humans), researchers may ascertain quantitative detail about the phenomenon in the model organism and try to explain these. Such pursuit of finer grain continues even within species; for example, different researchers focus, on the quantified details of circadian rhythms of young people or the aged, or on those with depression. There are many other ways researchers restrict the grain of phenomena. For example, they may focus on a specific component activity in the overall activity. As described above, circadian rhythms are generally characterized as endogenous generated oscillations that are entrainable to local conditions and are temperature compensated. Some circadian researchers focus on these aspects of circadian rhythmicity. Likewise, in memory research, investigators not only specialize on specific types of memory (e.g., episodic) but on encoding vs. storage or retrieval.

Once one focuses on more fine-grained phenomena, it becomes apparent that they have a host of different features. For example, endogenous rhythms in a particular activity exhibit a specific period, amplitude, and phase portrait. Explanatory research may be targeted at one such feature, or how multiple features interact. What is important to note is that what is put forward as an explanation for a specific feature may not be a mechanism. In the example discussed in section 3, researchers had not even advanced an empirically supported proposal about the mechanism when Winfree identified and proposed an explanation for a feature of the phenomenon. When a mechanism is known, it may contribute to the explanation of features of fine-grained phenomena, but even when that is the case, what is required for an explanation is not just the account of the mechanism but a demonstration of relations between variables (some or all of which may characterize features of components of a mechanism, such as their concentrations). In articulating the role of data graphs in such inquiries, Burnston (2016; see also Burnston, Sheredos, Abrahamsen, & Bechtel, 2014) refers to these relations between variables as *explanatory relations*. Individual research papers typically offer as explanations one or more empirically supported relations between variables.

Given the importance of articulating the features of a phenomenon for which explanations are sought, one might assume that these are delineated in advance of research on mechanisms. But, as Bechtel and Richardson (1993/2010) discussed, research on mechanisms sometimes leads to what they refer to as *reconstituting the phenomenon*—serious revisions in what the phenomenon is taken to be. There is another way, however, that work on mechanisms may provide features of fine-grained phenomena that become the targets of explanation. In the case of complex mechanisms, accounts such as those presented by Machamer, Darden, and Craver (2000), Craver (2007), Craver and Darden (2013), often fail to show how a mechanism is able to produce the phenomena. Bechtel and Abrahamsen (2010), Brigandt (2013), and (Baetu, 2015) have argued that in the case of mechanisms with non-sequential organization and non-linear operations, computational models are required to show how the mechanism generates the target phenomenon.³ Some

³ There has been substantial disagreement over whether computational models explain. Focusing on the mathematical model of the action potential advanced by Hodgkin and Huxley (1952), Weber (2008) defended it as explanatory while Craver (2008) argued that it did not explain since it did not describe the mechanism. Subsequently, Levy (2013) argued that Hodgkin and Huxley offered a deliberately abstract account but one that does explain the action potential in terms of component currents. The computational accounts discussed by Bechtel and Abrahamsen, Brigandt, and Baetu, in contrast, are tightly linked to mechanistic accounts—the differential equations in these models are drawn from the operations thought to constitute the mechanism. Although invoking mathematical derivations, these models are in the service of showing how mechanisms work and arguably in many cases one cannot show that the mechanism can produce the phenomenon except by using such models. At least in these cases, computational models seem to be critical to mechanistic explanation. Other mathematical models, such as those discussed by Chemero and Silberstein (2008), are models of phenomena, not mechanisms. The example from Winfree discussed below suggest a relatively clear way in which these models are explanatory as long as one is clear about what is being explained.

researchers find such models explanatory, but others contend from the computational model only shows that the mechanism can generate the phenomenon, but cannot explain how it does so. These researchers then treat the ability of the proposed mechanism to produce the target phenomenon as a further phenomenon requiring explanation.

In the following two sections I consider two cases in which research addressed a specific feature of a fine-grained phenomenon, and consider how the explanations researchers advanced differ from accounts of mechanisms.

2. Abstractly Representing a Phenomenon to Explain It

I turn in this section to a feature of the fine-grained phenomenon of circadian entrainment that was both predicted and shown to be a necessary feature of the phenomenon as a result of adopting an abstract representation of the phenomenon. In this case, the researcher both demonstrated and explained why the fine-grained phenomenon had a specific feature without reliance on an account of the responsible mechanism. Arthur Winfree, a mathematically inclined experimental biologist, performed this research during the period after the endogenous nature of circadian rhythms had been established but before the mechanism responsible for the general phenomenon had been proposed. During this period many circadian researchers focused their research on characterizing the more fine-grained phenomenon of circadian entrainment to light. The standard protocol for investigating circadian entrainment was to (a) establish the phase of endogenous oscillations by observing behavior while organisms were kept in darkness, (b) expose the organisms to light pulses of varying durations and intensities presented at different phases of the circadian cycle, and (c) determine how much the phase was advanced or delayed in each instance.

What researchers sought was not just the qualitative characterization of circadian oscillations being delayed or advanced depending on light exposure, but a detailed quantitative account of how rhythms were altered. To procure such an account of the phenomenon of entrainment, researchers needed to be able to identify a pattern in data they collected. The development of the phase response curve revealed such a pattern. Figure 2 shows a phase response curve from one of the first empirical studies of entrainment (DeCoursey, 1960). The time during the 24-hour period in which the light pulse is delivered is indicated on the x-axis with 0 indicating the time of activity onset for the two flying squirrels (*Glaucomys volans*) that were studied. The y-axis records the minutes by which the phase was advanced or delayed. This species is nocturnal so the 12 hours after the onset correspond to the period of expected darkness. What is clear is that light pulses during expected periods of light have no effect whereas those early in the expected dark period delay the phase of the rhythms and those late in the expected dark period advance it. (This makes sense since light pulses early in the dark period would correspond to light in the environment persisting longer than expected whereas those late in the dark period correspond to the light period starting earlier.)

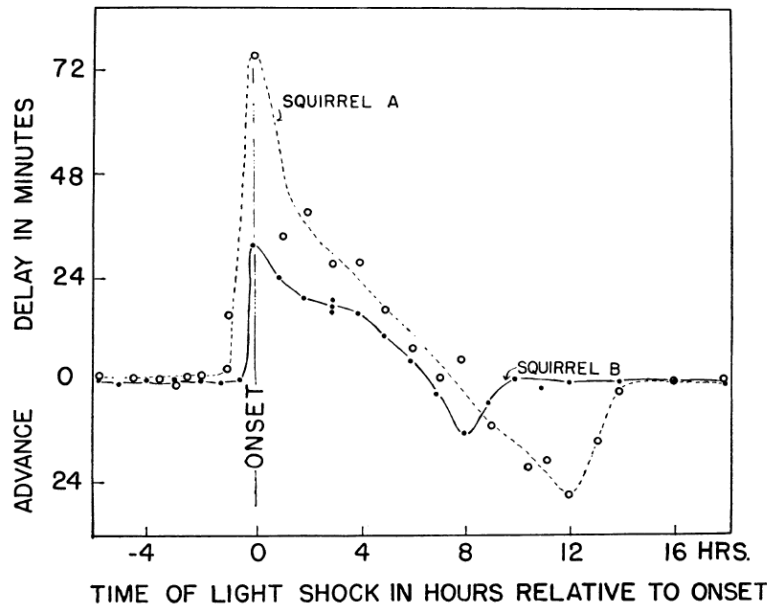


Figure 2. Phase response curve for two nocturnal flying squirrels from DeCoursey (1960).

Although phase response curves are widely employed to represent the phenomenon of circadian entrainment to light, they conceal an important distinction that is revealed using a different format, phase transition curves. As shown in Figure 3, these curves plot the old phase on the x-axis and the new phase on the y-axis. The dotted lines represent the case where the new phase exactly corresponds to the old phase (e.g., at the limit of a light pulse of 0 duration) and the dark wavy line shows how the new phase might vary from that. Figure 3A roughly corresponds to the pattern exhibited in Figure 2. When the light is given in the first part of the period (conventionally, phase 0 corresponds to the beginning of the expected light period) the phase is advanced. Light presented later in the period, on the other hand, delays the phase. (This graph is overly idealized. In fact during the light phase the solid line would closely match the dotted light and only show advance or delay during the dark phase.) Since the slope of the dotted line, which the actual curve approximates, is 1, this type of entrainment is referred to as Type 1. Essentially, this reflects the fact that during each period of 24 hours, circadian rhythms advance one period and the light perturbations do not change that.




Figure 3. Two types of entrainment shown using a phase transition curve. A. Type 1 B. Type 0. Figure from Lakin-Thomas, Coté, and Brody (1990), Figure 2.

However, in other cases the phase transition curve more closely fits the curve in Figure 2B, as Hastings and Sweeney (1958) found in their studies of *Gonyaulax polyedra*, a single-celled free-swimming plant that emits a faint blue light at night under circadian control. In their study, the organism shifted to a new phase based on the time of the pulse, regardless of the phase at which a light stimulus was presented. Sometimes the actual new phase was slightly delayed, sometimes advanced, with respect to the phase of stimulation, but overall it approximated a line with 0 slope. Accordingly, this is known as Type 0 entrainment.

Initially it appeared that the two types of entrainment, Type 1 and Type 0, occurred in different species. However, Winfree (1970) determined that both can occur in the same species in response to shorter (or weaker) versus longer (or stronger) entrainment stimuli; shorter stimuli yield Type 1 entrainment whereas longer stimuli yielded Type 0 entrainment. Winfree performed his studies using what was at the time one of the best-studied exemplars of circadian behavior, eclosion from the pupa into the mature fly in the fruit fly *Drosophila pseudoobscura*. In the wild eclosion occurs in the early morning hours and Pittendrigh had established that even when light, temperature, and other Zeitgebers are removed, eclosion time remains highly constrained, albeit advancing slightly each day since the free-running period of flies is somewhat shorter than 24 hours. Most studies of entrainment in *D. pseudoobscura* found Type 1 entrainment, but Winfree found that with extended light pulses, Type 0 entrainment occurred. Peterson (1980) subsequently discovered that both types of entrainment also occur in mosquitos.

Since the same organism shifts between Type 1 and Type 0 entrainment with longer (or more intense) light pulses, one might expect the transition between them would be smooth. Winfree demonstrated through a geometric argument that this was not possible. One can already see the problem in the phase transition curves shown in Figure 2. In Type 1 entrainment the phase transition curve ends up after 24 hours still aligned with the line with slope 1. With somewhat longer stimuli, the curve would depart further from that line at intermediate phases, but no matter how much it does so, it will return to the same relation to the line 24 hours later. It will never be transformed into the line for Type 0, which once it departs from the line with slope 1 never rejoins it. Winfree developed a perspicuous way of representing the lack of a smooth transition using a torus. The Cartesian representations in Figure 3 can be transformed into a torus by rolling the bottom edge backwards until it joins the top and the right edge backward until it joins the left edge. As illustrated in Figure 4, Type 1 entrainment involves a trajectory that passes once through the hole of the torus (providing another reason to call it Type 1), while Type 0 entrainment involves a trajectory that goes around the torus but never through the hole. It should be immediately clear that there is no way to transform a line going through the hole into one going around the outside without breaking and rejoining the line. This elegant geometrical representation reveals a mathematical truth: there can be no smooth transition from Type 1 to Type 0 entrainment. Rather, Winfree contended there will be an abrupt transition from Type 1 to Type 0 entrainment at a singularity, a point which conforms to neither Type 1 nor Type 0 entrainment.

A

B




Figure 4. A. Represented on a torus, Type 1 resetting corresponds to a line through the hole. B. Type 0 resetting involves an orbit around the torus but not through the hole.

Winfree also employed another representational format to demonstrate the absence of a smooth transition between Type 1 and Type 0 entrainment, one that also suggests what will happen at the singularity. Instead of a Cartesian coordinate system, a radial phase plot uses polar coordinates: different phases are plotted around the origin and the distances from the origin represent the values of a variable (e.g., amplitude) at those phases. In a radial phase plot the trajectory of an oscillatory processes appears as a closed figure known as the *limit cycle*. In Figure 5A the projecting spokes represent the phases of a 24-hour period and the circle represents the path through phase space of the system during each oscillation (it traverses each phase). The spokes need not be straight lines as shown here, but they may not cross. The spokes are referred to as isochrons since all the points along them represent the same phase. The radius of the circle represents the amplitude of the oscillation (concentric circles around the origin would represent oscillations of different amplitude). Note that at the origin the amplitude of the oscillation is reduced to 0; with no variation, no phase is defined.




Figure 5. A. Representation of a limit cycle in a radial phase plot with the different phases represented by lines emanating from the center and indicated by an hour. B. Type 1 entrainment shown in phase space with vector lines indicating how each point on the limit cycle is perturbed. C. Type 0 entrainment. Panels B and C from (Johnson, 1999).

The origin in the radial phase plot corresponds to what Winfree characterized as the singularity where neither Type 1 nor Type 0 entrainment applies. To see this, consider

how the phase shifts corresponding to entrainment can be represented in these radial phase plots. The arrows in panels B and C represent how each point on the limit cycle is perturbed. They either speed up the oscillation or slow it down. This is reflected in the fact that some of the arrows in panel B project to a later isochron, others to an earlier one. After a transient period the oscillation will return to a *limit cycle*, but with the new phase (this is not shown). Panel B represents Type 1 entrainment. What is important to note is that the origin remains inside the closed figure represented by the dashed line. In Type 1 entrainment the clock can be perturbed to any new phase but no stimulus sends it to the origin. Panel C, in contrast, represents Type 0 entrainment. Note that the arrows are mostly longer, representing a greater perturbation. Moreover, with such a stimulus, the clock cannot be set to any new phase but only those that fall between the isochrons 1 and 10. But more importantly, the origin is no longer within the closed figure defined by the dashed line. From Figure 5 one can determine that there is no transition from Type 1 to Type 0 entrainment that does not cause the closed figure to cross through the origin. Assume a gradual increase in the stimulus that generates Figure B. At some point the dashed line of the closed figure will cross the origin. As noted above, this is a singularity—a point that has no amplitude and hence no phase. What this represents is that whenever a stimulus moves the closed figure representing the oscillator to the origin, the amplitude of the oscillatory process drops to 0. With no amplitude, oscillation ceases.

As I discuss below, Winfree went beyond these geometric demonstrations to conduct an empirical investigation to determine exactly where the singularity occurred and that circadian oscillations ceased when organisms reached the singularity. But what is important to note is that he did not discover the existence of the singularity nor demonstrate that it had to occur through this further empirical inquiry. The occurrence of both Type 1 and Type 0 entrainment in the same species was discovered empirically, but the claim that when both types of entrainment occur in the same organism, there must be a singularity was demonstrated mathematically by abstractly representing the trajectories of the two types of entrainment, not through this further empirical inquiry. It is a mathematical truth (Lange, 2012) that explains that a singularity occurs.

In his further empirical investigations, Winfree made use of another discovery made by Pittendrigh. Pittendrigh (1966) found that keeping flies in constant light causes circadian rhythms to cease. These rhythms would start again with exposure to darkness, with the phase determined by the time of exposure to darkness. This provided an experimental protocol for precisely setting the phase of the circadian oscillators. Winfree created a two-dimensional array of dishes of pupae for which he could control their light exposure. He began by exposing the flies to constant dim blue light and then progressively blocked the light to different dishes over a three day period, beginning with those on the right (designated East in Figure 6a). This started oscillations in each population at the time light to it was blocked. After all the populations were in darkness and expected to be oscillating, he exposed them to pulses of light of varying duration, with those along the top (North in Figure 6a) getting the longest exposure. He then recorded when each fly emerged from its pupa. Since the initial exposure to darkness occurred over three days, there should be three populations, each exhibiting a full range of times of eclosion. Figure 5b shows this expected result. In each of the three populations, those first exposed to darkness (D) should eclose

first, with a wave moving from D to A. A light pulse experienced early in the subjective night should delay eclosion (B), whereas one late in the night should advance it (C). This should result in the oscillation of eclosion in each population from D to A, to B, to C, and again to D.

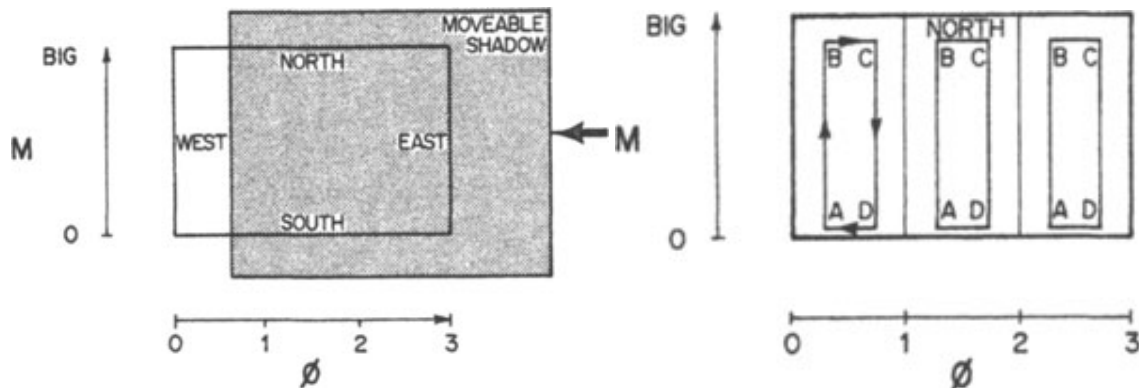


Figure 6. A. Winfree's experimental procedure of placing successive populations in darkness as a screen passed over their dishes over a three-day period. They were then exposed to a pulse of light originating at the top (North). This led to an expected pattern of eclosion involving three populations each exhibiting the cycle from D to C.

The dataset, consisting of the endogenous circadian time, length of light pulse, and eclosion time for each fly, was huge. Moreover, Winfree found it challenging to represent it since this required exhibiting both the endogenous and resulting phase and the duration of the stimulus. Initially he resorted to constructing a physical three-dimensional graph in which each wire represents an experiment at a given time (T) after transfer to darkness and log of exposure duration (S). Buttons on the wires represented the centroid of the times of emergence peaks for the corresponding group of flies (photographed in Figure 7a). Although he was only able to show the first third of the data before the structure became too complex to work with, the existence of a spiral pattern around an axis is already clearly apparent. The axis represents the singularity. Winfree comments: "There are no centroid data shown at this rotation axis, because following this perturbation, phase-resetting is erratic and flies emerge not in discrete peaks, but at all hours of the day, as is discussed below" (p. 331).

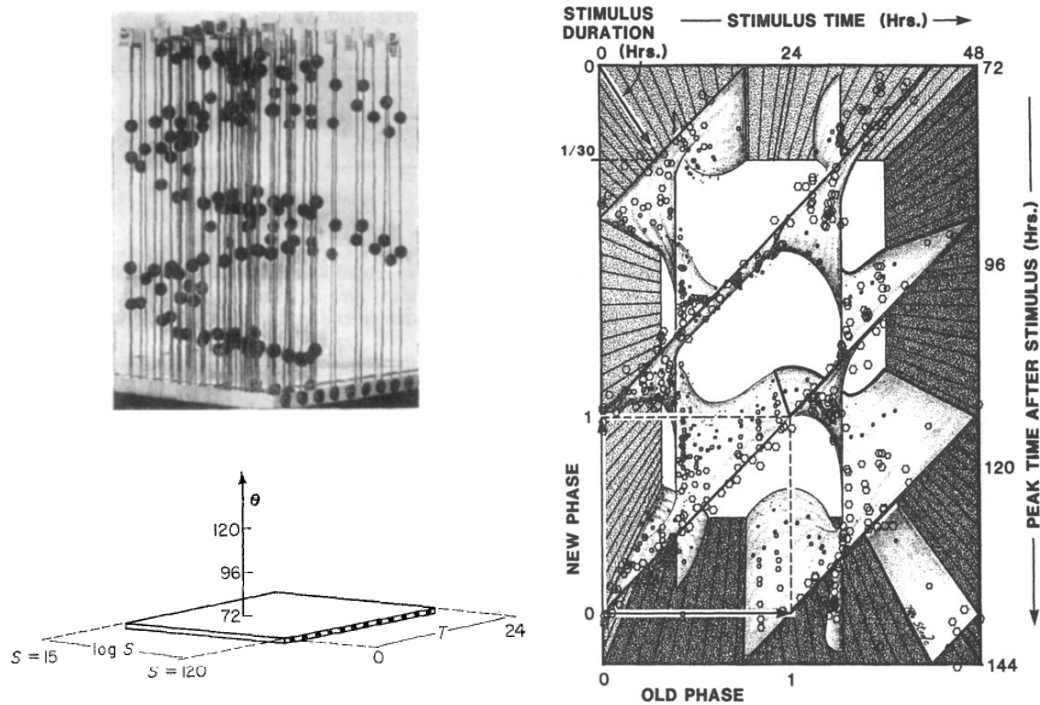


Figure 7. A. Photograph of a physical three-dimensional graph from Winfree (1970). B. Time crystal from Winfree (1980, p. 54).

Already in his 1970 paper Winfree developed a computer algorithm to define a surface through the eclosion points, which he termed a “resetting surface.” He wasn’t able to present it visually but described it as a “vertical corkscrew linking together tilted planes.” He further noted: “a corkscrew surface has a singularity, a central axis along which the slope is infinite.” He calculated that in his data the axis corresponded to a stimulus of 50 seconds at 6.8 hours after exposure to darkness. In *Geometry of Biological Time* (1980) and later work he presented this resetting surface in what he called a “time crystal” (Figure 7b). In the time crystal the x - and y -axes represent old and new phase respectively while the z -axis represents the duration of the stimulus. (Two complete periods of 24-hours are shown on the x -axis and three on the y -axis.) Each circle constitutes an eclosion event. Winfree adopted the scheme of representing responses to shorter stimuli (represented in the foreground) with larger circles and response to longer stimuli with smaller circles. As a result, one sees Type 1 resetting in the foreground and Type 0 in the background. The plane defined by Winfree’s algorithm appears as a surface that wraps about the singularity. The singularity, which is actually shown twice, once in the first period of 24-hours and again in the second 24-hour period, is the perpendicular axis around which the surface turns in the fashion of a circular staircase.

To determine what would happen at or near the singularity, Winfree performed 44 additional experiments using stimuli that were very near to where he calculated the singularity was to be found. As the stimulus approached the singularity, the eclosion pattern ceased to show a clear rhythm—the peaks of the distribution broaden significantly and even became indistinguishable. When an oscillation was suggested, its period often

varied from 24 hours. Winfree concluded that the flies had become arrhythmic and the circadian clock had essentially stopped. What he meant by this is that although the physiological processes that constitute the clock are still occurring, they no longer result in oscillation; it was the phase of the oscillation that represented time, and there no longer was a phase.⁴ The clock would remain stopped until a new light stimulus was presented, at which time oscillations would resume at the phase it was at when it was interrupted.

Winfree's experimental studies allowed him to identify the time and stimulus strength of the singularity and demonstrate what happened in response to a stimulus corresponding to the singularity. Nonetheless, Winfree had not only predicted the occurrence of the singularity but also explained why there had to be one before these experiments. Using abstract representations, such as the torus and the phase diagrams, he showed that the singularity was a necessary feature of any system exhibiting both Type 1 and Type 0 entrainment. Moreover, this explanation was independent of knowing the mechanism of entrainment. Subsequent to Winfree's analysis, researchers have identified key components of the mechanism of entrainment, but we still do not know why oscillation stops in response to a stimulus corresponding to the singularity. Nonetheless, Winfree provided an explanation of why there is and must be a singularity in the entrainment of organisms that exhibit both Type 1 and Type 0 entrainment. This particular feature of the phenomenon entrainment that occurs in at least some species is explained and its explanation does not depend on knowing the responsible mechanism but on understanding the mathematical fact that there must be a discontinuity or singularity between the range of stimuli that generate Type 1 and the range that generates Type 0 entrainment.

Some may question whether what Winfree's explanation of the singularity should count as an explanation. It is certainly not a causal or mechanistic explanation, and if one holds that it is causes (Salmon, 1984, 1998) or mechanisms (Craver, 2007) that explain, then they will not count what Winfree offered as an explanation. Salmon and Craver, however, were arguing against an alternative account of explanation. On the deductive-nomological (D-N) account of explanation (Hempel, 1965, 1966) one explains an event by showing that its occurrence followed deductively from laws and initial conditions. Winfree's account can be seen as a deduction—from the conjunction of the two types of entrainment, he showed that a singularity must occur. But while the characterizations of the two types of entrainment may be represented as empirical generalizations, they are not what would generally be regarded as candidate laws. What does the explanatory work, however, is not the accounts of the two types of entrainment but the mathematical demonstration that no smooth transition is possible between them.

While Winfree's explanation might fall under the D-N account, it leaves open the question of when one should seek a mechanistic explanation and when a D-N style explanation might suffice. Appealing to yet another philosophical account of explanation may suggest a framework for address this. Bromberger (1966, 1968) and (van Fraassen, 1980) treat

⁴ With the discovery that it was the concentrations of proteins such as PER that oscillated, the clock stopping can be understood as the concentration of these proteins reaching a constant level and no longer oscillating.

explanations as answers to questions about why something happens. For van Fraassen, explanation is pragmatic and relative to the context in which the question is asked and what is required to answer it is an account of why some condition arose rather than some other. On such a pragmatic understanding of explanation, Winfree is offering an explanation since he is answering the question why there must be a singularity. In this context, a D-N style account provides the answer to the question posed. That Winfree views his account as an explanation is clear in (Winfree, 1987) when he extends his account to show why, given tidal patterns in different locations in an ocean, there must be a singularity at which there are no tides. Although he is not offering a causal or mechanistic explanation, or even a D-N explanation, Winfree is answering a question about why something occurs.

Those who don't see Winfree as offering an explanation might argue that what he offered was only a discovery strategy: he predicted the singularity and his subsequent research showed that it occurred as he predicted. In this case the specific feature of the phenomenon of entrainment was not known before Winfree performed his analysis—he predicted it based on his analysis. On the D-N account, prediction and explanation were viewed as linked—the same argument could serve to predict and explain. Craver's criticism of the D-N model focused on this feature. Using Aristotle's example of the flagpole, Craver argued that one could predict the height of the flagpole from the length of its shadow, but that did not explain it. While it is certainly the case that sometimes the reasoning that supports a prediction does not suffice for explanation, in fact from a sufficiently detailed explanation, even an account of a mechanism, one can derive predictions. In fact, many tests of mechanistic hypotheses depend upon making predictions about how a proposed mechanism will behave and determining whether the actual system behaves in that way. While not all predictions rely on having a correct explanation, explanations often facilitate predictions and, as in this case, the discovery of a new feature of a phenomenon. Once the feature was established empirically, the very reasoning that led to its discovery also suffices to explain its occurrence.

In his 1987 book *The Timing of Biological Clocks*, a *Scientific American* book in which he presented his account of why circadian entrainment in many species exhibits a singularity, Winfree commented "How does phase resetting come about? Its results can be described without describing the process—a piece of good fortune, since no one yet knows the mechanism of a single circadian clock." Subsequently researchers have learned a good deal about the circadian clock mechanism in many species and even some of the details about it can be entrained by light stimuli. Filling in the account of the mechanism involved in entrainment will not supplant the explanation Winfree offered of the singularity. It will, though, address a new explanatory challenge to which Winfree's discovery gave rise: explaining how a stimulus corresponding to the singularity actually stops the clock mechanism by stopping the oscillations of circadian proteins. This challenge is directed at a different fine-grained phenomenon than Winfree addressed—the response of the mechanism itself to the stimulus that Winfree had shown would stop circadian rhythmicity. Note that the mechanism itself figures in this fine-grained phenomenon whereas it did not in the feature of the phenomenon that Winfree discovered and for which he offered an explanation. This highlights the fact that some fine-grained phenomena are in

fact characterized in terms of a mechanism and the challenge is to explain why the mechanism exhibits that feature of the phenomenon. This is illustrated in the next case.

3. Abstractly Representing a Mechanism to Explain Why It Produces the Phenomenon

Shortly after Winfree bemoaned the lack of knowledge of any circadian clock, Hardin et al. (1990) offered the first proposal for a TTFL mechanism. In section 1 I provide a verbal narration of how the initially proposed feedback loop would generate oscillations. This description does not establish whether the proposed mechanism would generate sustained oscillations or dampen over time. Those interested in this more specific feature turned to a computation model (Goldbeter, 1995) of the mechanism.⁵ In the subsequent decade many more parts were discovered and circadian researchers proposed mechanisms involving multiple feedback loops, both positive and negative (Figure 8 shows the conception of the circadian clock mechanism in mammals that was arrived at by 2005). With these discoveries, computational models became even more critical to determining how a proposed mechanism would behave. Computational models proposed by Leloup and Goldbeter (2003, 2008), (Gonze, 2011), and others provided support for the claim that a mechanism of this type could generate sustained oscillations.

⁵ Bechtel and Abrahamsen (2010, 2011) refer to mechanistic explanations that rely on computational modeling to establish that they exhibit specific dynamical behavior as *dynamic mechanistic explanations*.

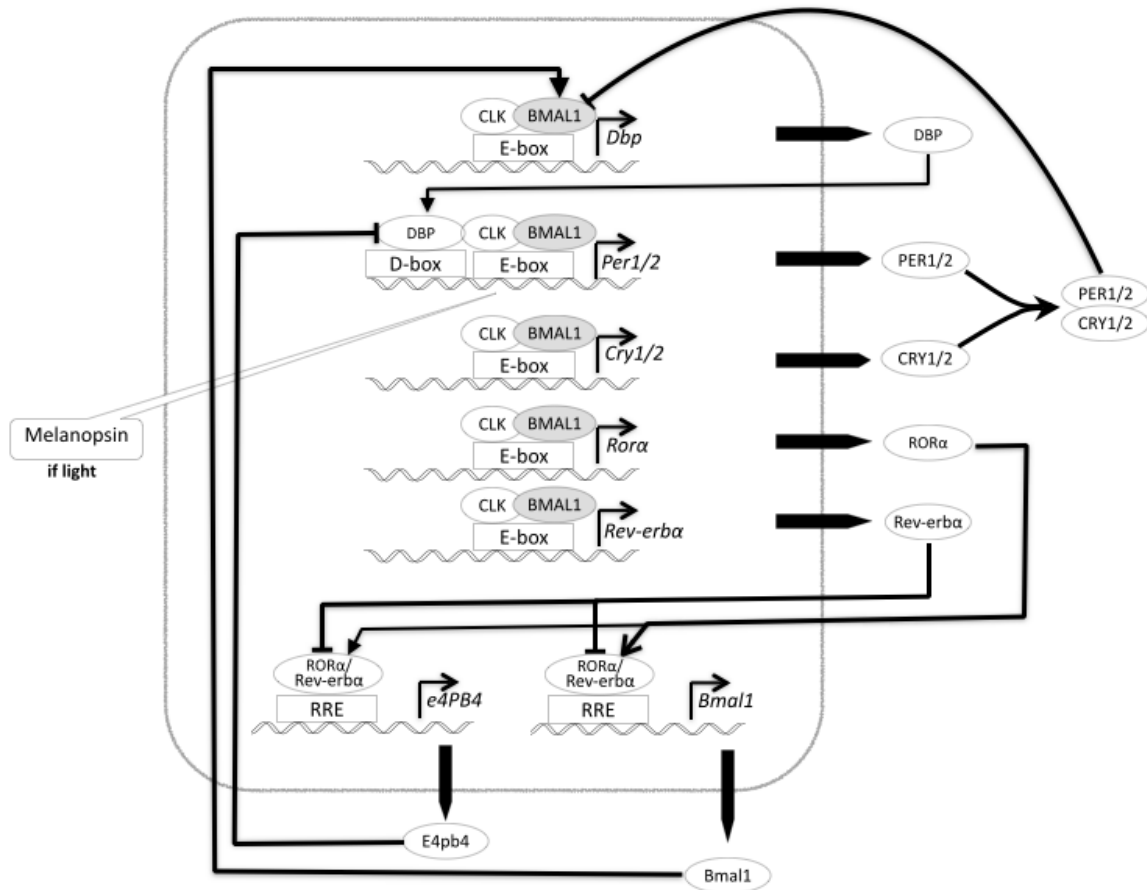


Figure 8. A representation of the major components in the mammalian circadian clock as understood circa 2005.

Computational models can show that a proposed mechanism is adequate to generate the specific feature of the phenomenon in question—sustained oscillation. But given the number of parts proposed, these models require dozens or hundreds of differential equations. While one can acquire an intuitive sense of why a single delayed negative feedback loop generates oscillations, circadian researchers lacked such an intuitive understanding of how these more complex mechanisms and models would behave and sought to understand why they generate sustained oscillations. The proposed mechanism is claimed to generate oscillations as a result of oscillations in the concentrations of components of the mechanism such as PER and BMAL1. These are proposed to oscillate in a circadian fashion and their doing so became the feature of the fine-grained phenomenon to be explained.

One approach to explaining this feature has been to focus on some part of the mechanism and view it as responsible for the oscillation in concentration of the other parts. Given that the mechanism involves several feedback loops, one might think one or another is what drives the oscillation of the whole mechanism. Researchers have not been able to pursue this line of inquiry through empirical experimentation, but some investigators experimented on computational models. Their strategy is to build a model that exhibits

circadian oscillations and then alter parts of the model in ways that correspond to removing or fixing the state of some of the parts of the mechanism. However, different modelers pursuing this strategy have arrived at diametrically opposite results (Smolen, Baxter, & Byrne, 2002; Relógio, Westermarck, Wallach, Schellenberg, Kramer, & Herzel, 2011). The differences in conclusions reflect differences in the details of how the models are constructed. Although this line of research has not yet generated definitive results, it may in the future. However, I will not discuss this approach further in this paper.

An alternative approach is not to look to a part of the mechanism to explain the phenomenon but to the organization of the whole. The challenge is that the overall organization is very complicated (Figure 8 in fact presents a simplified view) and it is not obvious why it would generate a sustained oscillation. Ueda and his collaborators developed a strategy for identifying what it is about this organization that is responsible for the feature of sustained oscillation (Ukai-Tadenuma, Kasukawa, & Ueda, 2008; Ukai-Tadenuma, Yamada, Xu, Ripperger, Liu, & Ueda, 2011; Hogenesch & Ueda, 2011). They began by re-representing the mechanism shown in Figure 8 using the scheme shown in Figure 9. This representation places at the center the three promoter boxes on the different genes in Figure 8: the E-box, D-box, and RRE. Since one or more of these boxes is present on each gene that is part of the circadian clock as well as on many other genes whose expression is controlled in a circadian fashion, they refer to these as *clock controlled elements (CCEs)*.

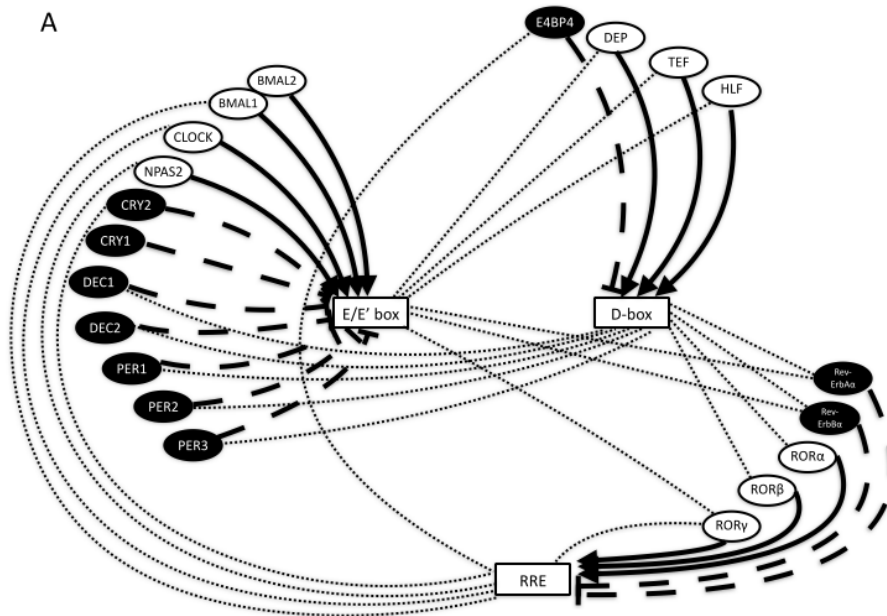


Figure 9. Ueda’s schema for representing the clock mechanism shown in Figure 8 in which the promoter boxes are made central and the gene/proteins serve to link activity between the various promoters.

Figure 9 downplays the processes of transcription and translation that are shown in Figure 8. First, the distinction between genes and proteins is collapsed. The CCEs regulate transcription and the dotted lines between CCEs and the ovals suggest this control over

transcription. But the fact that the names in the ovals are in capital letters and not italicized suggests they are proteins. This fits with the arrows and dashed, edge-ended lines that link the ovals back to the CCEs; these indicate the activity of proteins inhibiting or promoting transcription. But the distinction between genes and proteins and the operations in which each participate ultimately doesn't matter for the analysis Ueda is advancing. His strategy is to abstract from the details of the genes and proteins and simply treat them as intermediaries between the CCEs.

The foregrounding of the CCEs was largely motivated by Ueda's experimental finding that when he inserted destabilized luciferase genes into the region regulated by the CCEs in a cell culture system and recorded the timing of maximum bioluminescence, he found that each CCE was most active at a different time. Although the precise time varies by tissue, in the suprachiasmatic nucleus, thought to be the locus of the central clock in mammals, E-boxes are most active in the day, D-boxes about five hours later (evening) and RREs about eight hours later (night). This suggests that the activity of the different boxes plays a central role in generating a 24-hour oscillation. The genes/proteins are simply the means by which the CCEs affect each other.

To make the implications of this representation clearer, in a subsequent diagram (Figure 10A) Ueda no longer displayed the genes/proteins and instead inserted a single arrow or edge-ended line for all the pathways between a given CCE and each of the other two. Thus, the arrow between the E/E' box and the D-box indicates that one or more genes regulated by the E/E' box are synthesized into proteins that function to enhance expression of genes with a D-box. The two dotted, edge-ended lines indicate that the D-box and RRE have inhibitory effects on the period and amplitude of oscillation of genes with E/E' boxes.




Figure 10. A. A representation of the clock mechanism that abstracts from the genes and proteins and shows only when there is at least one pathway from a gene controlled by one box to the activation or inhibition of the box another gene. B. A decomposition of the figure in A into two motifs, a repressilator and a delayed negative feedback loop.

The goal of abstracting from all the genes and proteins and generating a skeletal representation of the system was to elicit if possible an explanation for why the mechanism generates oscillations within it. By removing the identity of the genes and proteins and focusing only on how they serve to connect the CCEs, Ueda is focusing his attention on the organization. A potent way to represent organization and abstract from the properties of

the components or the details of the operations the components perform is to develop a network diagram in which entities appear as nodes and any connections between entities as edges. Both Figures 9 and 10 are network diagrams, whereas Figure 8 is more properly thought of as a mechanism diagram. While Figure 10A is still complex and does not itself provide an explanation, Ueda recognized that this network could be decomposed into the two networks shown in Figure 10B. The one on the left is a repressilator and the one on the right a negative feedback loop. These correspond to what are elsewhere called *motifs*. Working at the level of the actual components in yeast gene and protein interaction networks, Alon and his collaborators (Milo, Shen-Orr, Itzkovitz, Kashtan, Chklovskii, & Alon, 2002; Alon, 2007a, 2007b) identified subnetworks of two to four nodes similar to those in Figure 10B. Their attention was drawn to these subnetworks, which they termed *motifs*, because they occurred frequently. By making minimal assumptions about the actual entities involved, they analyzed the behaviors each motif would support and proposed that was the contribution the specific subnetworks would make to the larger network. Ueda arrived at motifs in a different manner (by abstracting from the numerous genes and proteins through which CCEs affect each other), but the notion of motif applies here as well. Essentially, the two networks shown in Figure 10B are implemented many times in Figure 9 depending on which genes and proteins serve as intermediaries.

What renders the notion of motif powerful is that one can establish how, within a range of parameters, any subnetwork implementing the motif will behave. Indeed, both of the motifs Ueda extracted had already been demonstrated, both in computational models and in synthesized organisms, to be capable of generating sustained interactions. The delayed negative feedback motif played a fundamental role in theorizing about the clock mechanism from the 1960s and was incorporated into the TTFL model. Its origins lay much earlier. Since Ktesibios employed it in his water clock in the second century BCE, it has been rediscovered and employed to maintain features of a system (e.g., temperature in a room) at a constant level (Mayr, 1970). In the second quarter of the 20th century this use of negative feedback was celebrated by the cyberneticists (Wiener, 1948). But in a variety of applications researchers also recognized that many systems implementing negative feedback would not settle to a steady-state but would oscillate around the desired value. Beyond empirical observations of oscillation in negative feedback systems, numerous engineers developed mathematical analyses of feedback systems, of which one of the most influential was proposed by van der Pol (1920), an electrical engineer. These analyses showed that sustained oscillation was possible.

Although its prominence arose more recently, the repressilator was also recognized as a motif capable of generating oscillations. In the process of designing a synthetic oscillator they planned to incorporate into *E. coli*, Elowitz and Leibler (2000), constructed a computational model of a repressilator circuit and showed that, with appropriate parameter values, it generated sustained oscillations. They then inserted into bacteria genes that interacted in the manner indicated in the motif as well as a green fluorescent reporter. They then observed oscillations with a period of several hours.

Given that both motifs had been shown to generate oscillations, Ueda argued that the ability of the complex mechanism to generate sustained circadian rhythms is due to the fact

that the organization found in the complicated mechanism shown in Figure 9 realizes these two motifs. In Ueda's analysis, the motifs were identified not in individual circuits, but in an abstract representation of a complicated network. The arrows in the motifs represent multiple transcription/translation relations in which product proteins affect another promoter. Thus, instead of a local circuit, it is the network as a whole that implements the two motifs. Yet it is the motifs that explain why the circadian mechanism exhibits sustained oscillations. It is explained by the way the parts are organized and that their operations the parts perform relate parts in the manner reflected in the motif. As Levy and Bechtel (2013) discuss with respect to Alon's work, the analysis of motifs makes only minimal claims about the nature of the components corresponding to the nodes and the operations they perform on each other. Any system organized according to the motif in which these minimal conditions are met will exhibit the associated behavior. In this sense, motif analyses, and other network analyses, offer general accounts that apply to all instances in the manner proposed for D-N explanation. The motifs, however, are not laws, but principles of organization (design principles) that assert that any system implementing the organization will exhibit the specified behavior.

While abstraction facilitates explaining the feature of the generation of rhythms that Ueda was interested in, it can impair the ability to explain other features. A different feature of the phenomenon that Ueda did not focus on is that a circadian oscillation is extremely slow for chemical reactions, completing a cycle only once every 24 hours. This depends on delays within the mechanism. One of the important operations that Ueda's account abstracts from is that in order to get into the nucleus where they can act on a CCE, Per1 and Cry1 must function as a dimer to reenter the nucleus, with the formation of a dimer occurring over an extended period of time. Other features of the phenomenon require an intermediate level of abstraction. Considerable interest in recent years has been directed at the robustness of the oscillation to alteration of individual components in the mechanism. This is partly explained by the presence of multiple orthologs of key proteins (e.g., Cry1 and Cry2) in the mammalian clock. Knocking out just one has little effect since the other can compensate by increasing the amplitude of their oscillations and maintain the functioning of the same motif (Baggs, Price, DiTacchio, Panda, FitzGerald, & Hogenesch, 2009).

Circadian researchers devoted major efforts in the 1990s and early 2000s to developing a detailed mechanistic model of the mammalian circadian clock. As useful as this model is, it does not explain why the mechanism generates sustained circadian rhythms of the other features noted in the previous paragraph. According to the model, components of the mechanism themselves exhibit sustained circadian rhythms. This calls for explanation. By abstracting from the details of the genes and proteins involved, Ueda was able to demonstrate that the organization of components in the modeled mechanism realizes two motifs known to generate sustained oscillations. The abstract motif analysis revealed the design principle realized in the complicated mechanism and these are advanced as explanations of the ability of components in the mechanism to exhibit sustained oscillations.

4. Conclusions

A widely accepted view is that explanation in biology involves identifying the mechanism responsible for a phenomenon. Indeed, well supported mechanistic models have been advanced for many biological phenomena such as cell division, protein synthesis, and circadian rhythms. But a great deal of research in biology is directed not at such general phenomena but at specific features of far more fine-grained phenomena and often identifying the mechanism is not what explains these specific features.

To illustrate two of the strategies researchers invoke in explaining specific features of fine-grained phenomena, I have presented two cases involving research on circadian rhythms. While textbooks might characterize circadian rhythms in a general way, researchers typically focus on a particular feature of a more fine-grained phenomenon, such as entrainment by light. The first example focuses on one such feature, the occurrence of a singularity, which Winfree both predicted and explained through the same type of argument. Upon discovering that some organisms exhibit both Type 1 and Type 0 entrainment, Winfree demonstrated using a geometrical argument that there cannot be a smooth transition between the two types of entrainment. Rather, any transition from Type 1 and Type 0 entrainment has to go through a singularity at which the amplitude of the oscillation declines to 0. An entrainment stimulus that causes the clock to reach the singularity stops the clock. Winfree's explanation for the impossibility of smooth transition between types of entrainment did not depend on details about the mechanism. Once the mechanism of entrainment is understood, it can contribute to understanding how the system responds when it receives an input that drives it to the singularity. But the explanation for there being a singularity does not depend on the details of the mechanism.

The second case involves research that was carried out in the wake of the discovery of many of the parts and operations of the circadian clock. Merely identifying the parts, operations, and organization of the mechanism did not explain why the mechanism exhibits sustained oscillations. By abstracting from the details of the mechanism, Ueda was able to identify two motifs (negative feedback and the repressilator) that are realized in the organization of the complicated mechanism that was proposed for the clock. Mathematical and experimental investigations of these motifs demonstrate that they are design principles that result in sustained oscillations in systems in which they are implemented. While the mechanism is highly relevant to explaining the general phenomenon of circadian rhythms, the specific phenomenon of sustained oscillation within the mechanism is explained in terms of the design the design principles the mechanism implements in its organization.

To make this argument, I have had to make clear what specific feature of a fine-grained phenomenon is the target of a given explanation. What counts as explanatory depends critically on what a researcher is trying to explain. Following Bogen and Woodward (1988), most accounts of mechanistic explanation have construed phenomena quite broadly, treating, for example, the generation of circadian rhythms or the synthesis of proteins as single phenomena. In textbooks and sometimes in review articles scientists do speak this way, but in actual research and the resulting journal articles, phenomena are characterized far more narrowly and particular features are addressed. Winfree and others focusing on entrainment were not focused on the general fact that circadian rhythms can be entrained

to light conditions, but specific patterns of phase advances or delays in response to varying durations of light exposure. Ueda was not concerned with circadian rhythms in general, but with how the accepted account of the mechanism is able to generate sustained rhythms. Focusing only on the specific feature of the phenomenon for which explanation is sought helps to explain why only some information may be relevant to explaining it and why details of the mechanism may not advance the specific explanatory goals.

Finding the relevant explanatory principle to explain a specific feature of a phenomenon often requires developing the requisite abstract analysis. Had the details of the mechanism of entrainment been known, that would not have answered Winfree's explanatory quest. That required a mathematical analysis that showed why a singularity is required given the nature of the phenomenon. Likewise, in Ueda's case, answering the question of why the mechanism oscillated required identifying the underlying design principles and showing that systems implementing these design principles generate sustained oscillations. Developing explanations involves both identifying the specific features of a phenomenon for which an explanation is sought and developing the (often abstract) account that is tailored to explain those features.

A reason to focus on identifying the features of a phenomenon that are the target of particular explanations is that what is required to provide explanation will vary with the target. Just how much an account of the mechanism is required will vary. In arguing that often the mechanism is either not needed or insufficient to provide the explanation required, I am rejecting the claim that mechanistic accounts alone explain. In some cases, the explanations produced come closer to D-N explanations—researchers offer a general account and show what follows from it. What provides the general account might not be a law. In Winfree's case, it was a mathematical truth that could be shown to apply to the fine-grained phenomenon when it was characterized abstractly. In Ueda's case, it was two motifs that could be modeled mathematically to show how any system realizing them would behave. These are just two examples. Other cases will involve a variety of relations between variables that Burnston (2016) refers to as explanatory relations. The more general point is that to understand what provides the desired explanation in a given research endeavor depends on what feature of a fine-grained phenomenon a researcher seeks to explain.

References

- Alon, U. (2007a). *An introduction to systems biology: Design principles of biological circuits*. Boca Raton, FL: Chapman & Hall/CRC.
- Alon, U. (2007b). Network motifs: Theory and experimental approaches. *Nature Reviews Genetics*, 8, 450-461.
- Baetu, T. (2015). From Mechanisms to Mathematical Models and Back to Mechanisms: Quantitative Mechanistic Explanations. In P.-A. Braillard & C. Malaterre (Eds.), *Explanation in Biology. An Enquiry into the Diversity of Explanatory Patterns in the Life Sciences*. Dordrecht: Springer.
- Baggs, J. E., Price, T. S., DiTacchio, L., Panda, S., FitzGerald, G. A., & Hogenesch, J. B. (2009). Network features of the mammalian circadian clock. *PLoS Biol*, 7, e1000052.

- Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36, 421-441.
- Bechtel, W., & Abrahamsen, A. (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., & Abrahamsen, A. (2011). Complex biological mechanisms: Cyclic, oscillatory, and autonomous. In C. A. Hooker (Ed.), *Philosophy of complex systems. Handbook of the philosophy of science* (Vol. 10, pp. 257-285). New York: Elsevier.
- Bechtel, W., & Richardson, R. C. (1993/2010). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Cambridge, MA: MIT Press. 1993 edition published by Princeton University Press.
- Bogen, J., & Woodward, J. (1988). Saving the phenomena. *Philosophical Review*, 97, 303-352.
- Brigandt, I. (2013). Systems biology and the integration of mechanistic explanation and mathematical explanation. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 44, 477-492.
- Bromberger, S. (1966). Why-questions. In R. C. Colodny (Ed.), *Mind and cosmos: Essays in contemporary science and philosophy* (Vol. 68-111). Pittsburgh, PA: University of Pittsburgh Press.
- Bromberger, S. (1968). *An approach to explanation*. Paper presented at the In R. J. Butler, ed., *Analytic Philosophy: Second Series*. Oxford: Blackwell. Pp. 72-105.
- Burnston, D. C. (2016). Data graphs and mechanistic explanation. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 57, 1-12.
- Burnston, D. C., Sheredos, B., Abrahamsen, A., & Bechtel, W. (2014). Scientists' use of diagrams in developing mechanistic explanations: A case study from chronobiology. *Pragmatics and Cognition*, 22, 224-243.
- Chemero, A., & Silberstein, M. (2008). After the philosophy of mind: Replacing scholasticism with science. *Philosophy of Science*, 75, 1-27.
- Craver, C. F. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. New York: Oxford University Press.
- Craver, C. F. (2008). Physical law and mechanistic explanation in the Hodgkin and Huxley model of the action potential. *Philosophy of Science*, 75, 1022-1033.
- Craver, C. F., & Darden, L. (2013). *In search of mechanisms: Discoveries across the life sciences*. Chicago: University of Chicago Press.
- De Mairan, J.-J. d. O. (1729). Observation Botanique. *Histoire de l'Academie Royale Sciences*, 35.
- DeCoursey, P. J. (1960). Daily light sensitivity rhythm in a rodent. *Science*, 131, 33-35.
- Edmunds, L. N. (1988). *Cellular and molecular bases of biological clocks: Models and mechanisms for circadian timekeeping*. New York: Springer-Verlag.
- Elowitz, M. B., & Leibler, S. (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.
- Gonze, D. (2011). Modeling circadian clocks: From equations to oscillations. *Central European Journal of Biology*, 6, 699-711.

- Green, S. (2015). Revisiting generality in biology: systems biology and the quest for design principles. *Biology & Philosophy*, 30, 629-652.
- Halberg, F. (1959). Physiologic 24-hour periodicity: General and procedural considerations with reference to the adrenal cycle. *Zeitschrift für Vitamin-, Hormon- und Fermentforschung*, 10, 225-296.
- Hardin, P. E., Hall, J. C., & Rosbash, M. (1990). Feedback of the *Drosophila period* gene product on circadian cycling of its messenger RNA levels. *Nature*, 343, 536-540.
- Hastings, J. W., & Sweeney, B. M. (1958). A persistent diurnal rhythm of luminescence in *Gonyaulax polyedra*. *Biological Bulletin*, 115, 440-458.
- Hempel, C. G. (1965). Aspects of scientific explanation. In C. G. Hempel (Ed.), *Aspects of scientific explanation and other essays in the philosophy of science* (pp. 331-496). New York: Macmillan.
- Hempel, C. G. (1966). *Philosophy of natural science*. Englewood Cliffs, NJ: Prentice-Hall.
- Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to the conduction and excitation of nerve. *Journal of Physiology*, 117, 500-544.
- Hogenesch, J. B., & Ueda, H. R. (2011). Understanding systems-level properties: timely stories from the study of clocks. *Nature Reviews Genetics*, 12, 407-416.
- Johnson, C. H. (1999). Forty years of PRCs-What have we learned? *Chronobiology International*, 16, 711-743.
- Konopka, R. J., & Benzer, S. (1971). Clock mutants of *Drosophila melanogaster*. *Proceedings of the National Academy of Sciences (USA)*, 89, 2112-2116.
- Lakin-Thomas, P. L., Coté, G. G., & Brody, S. (1990). Circadian Rhythms in *Neurospora crassa*: Biochemistry and Genetics. *Critical Reviews in Microbiology*, 17, 365 - 416.
- Lange, M. (2012). What makes a scientific explanation distinctively mathematical? *The British Journal for the Philosophy of Science*.
- Leloup, J.-C., & Goldbeter, A. (2003). Toward a detailed computational model for the mammalian circadian clock. *Proceedings of the National Academy of Sciences*, 100, 7051-7056.
- Leloup, J.-C., & Goldbeter, A. (2008). Modeling the circadian clock: From molecular mechanism to physiological disorders. *BioEssays*, 30, 590-600.
- Levy, A. (2013). What was Hodgkin and Huxley's Achievement? *The British Journal for the Philosophy of Science*.
- Levy, A., & Bechtel, W. (2013). Abstraction and the organization of mechanisms. *Philosophy of Science*, 80, 241-261.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67, 1-25.
- Mayr, O. (1970). *The origins of feedback control*. Cambridge, MA: MIT Press.
- Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D., & Alon, U. (2002). Network Motifs: Simple Building Blocks of Complex Networks. *Science*, 298, 824-827.
- Peterson, E. L. (1980). Phase-resetting a mosquito circadian oscillator. *Journal of comparative physiology*, 138, 201-211.
- Pittendrigh, C. S. (1966). The circadian oscillation in *Drosophila pseudoobscura* pupae: A model for the photoperiodic clock. *Zeitschrift für Pflanzenphysiologie*, 54, 275-307.

- Relógio, A., Westermark, P. O., Wallach, T., Schellenberg, K., Kramer, A., & Herzog, H. (2011). Tuning the mammalian circadian clock: Robust synergy of two loops. *PLoS Computational Biology*, 7, e1002309.
- Salmon, W. C. (1984). *Scientific explanation and the causal structure of the world*. Princeton, N.J.: Princeton University Press.
- Salmon, W. C. (1998). *Causality and explanation*. Oxford: Oxford University Press.
- Smolen, P., Baxter, D. A., & Byrne, J. H. (2002). A Reduced Model Clarifies the Role of Feedback Loops and Time Delays in the *Drosophila* Circadian Oscillator. *Biophysical Journal*, 83, 2349-2359.
- Ukai-Tadenuma, M., Kasukawa, T., & Ueda, H. R. (2008). Proof-by-synthesis of the transcriptional logic of mammalian circadian clocks. *Nat Cell Biol*, 10, 1154-1163.
- Ukai-Tadenuma, M., Yamada, R. G., Xu, H., Ripperger, J. A., Liu, A. C., & Ueda, H. R. (2011). Delay in feedback repression by cryptochrome 1 is required for circadian clock function. *Cell*, 144, 268-281.
- van der Pol, B. (1920). A theory of the amplitude of free and forced triode vibrations. *Radio Review*, 1, 701-710, 754-762.
- van Fraassen, B. C. (1980). *The scientific image*. Oxford: Clarendon Press.
- Weber, M. (2008). Causes without mechanisms: Experimental regularities, physical laws, and neuroscientific explanation. *Philosophy of Science*, 75, 995-1007.
- Wiener, N. (1948). *Cybernetics: Or, control and communication in the animal and the machine*. New York: Wiley.
- Winfree, A. T. (1970). Integrated view of resetting a circadian clock. *Journal of Theoretical Biology*, 28, 327-374.
- Winfree, A. T. (1980). *The geometry of biological time*. New York: Springer Verlag.
- Winfree, A. T. (1987). *The timing of biological clocks*. New York: W.H. Freeman.