Scientists' Use of Diagrams in Developing Mechanistic Explanations: A Case Study from Chronobiology

Daniel C. Burnston[a,b], Benjamin Sheredos[a,b], Adele Abrahamsen[b,c], and William Bechtel[a,b]

[a] Department of Philosophy, University of California, San Diego
[b] Center for Circadian Biology, University of California, San Diego
[c] Center for Research in Language, University of California, San Diego

Abstract

We explore the crucial role of diagrams in scientific reasoning, especially reasoning directed at developing mechanistic explanations of biological phenomena. We offer a case study focusing on one research project that resulted in a published paper advancing a new understanding of the mechanism by which the central circadian oscillator in *Synechococcus elongatus* controls gene expression. By examining how the diagrams prepared for the paper developed over the course of multiple drafts, we show how the process of generating a new explanation vitally involved the development and integration of multiple versions of different types of diagrams, and how reasoning about the mechanism proceeded in tandem with the development of the diagrams used to represent it.

1. Introduction

An understanding of the importance of diagrams and other visual representations in reasoning has been pursued in at least two fields. First, cognitive scientists have found experimental support for the hypothesis that visual representations can play a potent *facilitative* role in human reasoning (Tversky, 2011; Hegarty, 2011). Second, philosophers of science, including the new mechanists, have posited an *explanatory* role for diagrams in science, particularly in biology (Perini, 2005; Sheredos, Burnston, Abrahamsen, & Bechtel, 2013; see Woody, 2000, for related discussions about diagrams in chemistry). Taken together, the two perspectives point to diagrams as important tools for reasoning in active science. How, then, might this role for diagrams be studied by philosophers of science?

In examining scientists' diagrams, philosophers of science usually have focused on particular interim or final products of this process—those published in books or articles

(e.g., Gooding's 2010 analysis of diagrams of arthropod anatomy and Nersessian's 2008 analysis of diagrams by the 19th-century physicists Faraday and Maxwell). However, scientists use visual representations as tools for reasoning throughout a nonpublic process of analyzing data and constructing explanations that can extend over months and years. Some highlights of this process may be captured by following published diagrams through time (what Sheredos, 2013, calls the study of lineages). However, the working versions that are missing from this record likely provide a better window on the scientist's active reasoning, since published versions may be altered to serve communicative ends. Thus, there is considerable value in the quite different ethnographic approach undertaken in a few interdisciplinary investigations of science (Alač, 2011; Osbeck, Nersessian, Malone, & Newstetter, 2010). Embedded in a laboratory for an extended period, such investigators potentially might examine in detail how scientists rely on a succession of working diagrams in performing the sophisticated reasoning involved in generating explanations over time. Unfortunately, this potential has not yet been fully realized (but see MacLeod & Nersessian, 2013, who included such diagrams in their ethnographic study of two systems biology laboratories).

In this paper we introduce a third, more efficient method of investigation in which the process by which scientists reason towards an explanation is traced by examining unpublished as well as published diagrams and associated manuscript drafts.[1] Researchers in a molecular chronobiology laboratory granted us access to such documents for an investigation that resulted in one journal article (Paddock, Boyd, Adin, & Golden, 2013). We already had ongoing interaction with them outside their laboratory, most regularly in a journal club involving researchers from several chronobiology laboratories. We maintained those contacts but did not substantially change or increase them after setting up this project, as our focus was on the unpublished manuscript drafts and diagrams. These were produced between October 2012 and June 2013, when the article was submitted. In February 2013 the first author (Paddock) began pursuing the implications of a finding not previously emphasized, which led to major reconstrual of an accepted mechanistic account in March-April 2013. We met with the first and fourth authors (project scientist and laboratory director, respectively) on March 4, 2013, to discuss the role of diagrams in their work and to arrange the case study involving their unpublished materials from this investigation in particular. We then waited to receive the diagram files and manuscript drafts until after the final version of their paper had been accepted for publication on August 13, 2013.

Our philosophical work on mechanistic explanation in science provided the framework for our examination of these materials. Although long neglected in 20th century philosophy of science, the importance of mechanistic explanation became apparent as attention extended beyond physics to biology. Having selected a phenomenon to investigate (e.g., some aspect of metabolism or respiration) life scientists typically attempt to (a) identify the particular mechanism most likely to be responsible at an appropriate level of organization (e.g.,

[1] This is not completely novel. Philosophers such as Nersessian have sometimes included diagrams from unpublished laboratory notebooks in their investigations.

organs, cells, molecules); (b) decompose this mechanism into its component parts and the operations they perform; and (c) show how, when appropriately organized and situated in an environment, these component parts and operations together can produce the phenomenon of interest (Bechtel & Richardson, 1993/2010; Bechtel & Abrahamsen, 2005; Machamer, Darden, & Craver, 2000). More recently, Bechtel and Abrahamsen (2012) have explored how these *basic mechanistic explanations* increasingly are being incorporated by life scientists into *dynamic mechanistic explanations* that add attention to the time course of the mechanism's operations. Ideally this involves computational modeling, our original focus, but other methods also can be of value, for example, running an animation, using the tools of dynamical systems theory to identify and display limit cycles, or plotting the time course of changes in a variable that can be linked to certain parts or operations.

In this paper, we coordinate mechanistic explanations with a type of representation that has received little attention in philosophy of science despite its importance in scientific practice: what we call *explanatory relations*. These are quantitative relationships, typically compared across experimental conditions or types of entities, that are important for explaining biological phenomena. They provide quantitative detail, often involving time, that is not present in a basic mechanistic account but usually linkable to it. Explanatory relations can be represented in data graphs, and scientists may devote considerable time and energy to manipulating and reorganizing their data to find the most informative explanatory relations. (For a more extensive discussion of data graphs and explanatory relations, see Burnston, forthcoming.)

This framework suggests several epistemically distinct types of diagrams. We focus here on two of these. *Explanatory relations diagrams* display data from two or more experimental conditions, often as line graphs but also in other formats for displaying quantitative relationships. *Mechanism diagrams* characterize the mechanism proposed as responsible for a phenomenon by providing a two-dimensional spatial representation of its parts, the operations they perform, and how they are organized spatially, temporally, and/or causally. Our study suggests that explanatory relations and mechanism diagrams mutually constrain each other, and that each is important for explanation. Because they offer different resources for reasoning and afford distinct kinds of inferences, the interaction between them is important for understanding the mechanism. (See Sheredos et al., 2013, for a more extensive discussion of diagrams and cognitive affordances.)

We set the stage for our case study in section 2, where we introduce the phenomenon of circadian rhythmicity in cyanobacteria and present the mechanism generally accepted in 2012, when the research was performed. In sections 3 and 4 we describe some of Paddock et al.'s data, the explanatory relations they extracted, and the reasoning that led them to question two major aspects of the previously accepted underlying mechanism. In seeking to resolve these questions, they developed a series of explanatory relations diagrams from which they inferred revisions to their mechanism diagrams. Their published article included only the final versions of these diagrams, most notably a single figure that represented an important conclusion by displaying the relevant explanatory relations diagrams and mechanism diagrams side by side.

2. Explaining the Phenomenon of Circadian Rhythmicity in Cyanobacteria

In chronobiology, the primary phenomenon to be explained is circadian rhythmicity: regularly recurring oscillations in activity with a period of approximately 24 hours, generated endogenously (internally) but entrainable to such environmental cues as the earth's daily light-dark cycle. Circadian rhythms have been observed in a host of physiological and behavioral activities in humans and other mammals (e.g., sleep, immune responses, and metabolism). They also occur, however, in a wide range of other organisms, including fruitflies (e.g., eclosion), plants (e.g., leaf movement), and even in cyanobacteria (e.g., photosynthesis, regulation of gene expression).

Research seeking explanations of circadian rhythms like these have focused predominantly, though not exclusively, on circadian rhythmicity at the intracellular level. Within many cells is a *molecular clock* mechanism—a system of molecules that undergo regular cycles in their causal interactions approximately every 24 hours. It is adaptable to, but not dependent upon, external cues. Much of the success of chronobiology has come in identifying and decomposing the molecular clock mechanism in a variety of organisms (see Dunlap, Loros, & DeCoursey, 2004). Paddock et al. focused on circadian rhythmicity of gene expression in a single-celled cyanobacterium, *Synechococcus elongatus*. Their published paper (Paddock, Boyd, Adin, & Golden, 2013) offered new data, new diagrams, and—mediated by a series of unpublished diagrams—a reconstrual of a key mechanism.

The central clock in *S. elongatus* is a posttranslational mechanism comprising three proteins—KaiA, KaiB, and KaiC—and the key oscillation involves the phosphorylation and dephosphorylation of KaiC at two sites, T432 ('T' for threonine) and S431 ('S' for serine).[2] At dawn KaiC is not phosphorylated, and during a 24-hour period it cycles through three phosphorylation states before arriving back at the original state. First the T site is phosphorylated, then both the S and T sites, then only the S site, then back to none. The reactions changing the state of KaiC are mediated early in the day by KaiA, which promotes phosphorylation at both sites, and later in the day by KaiB, which inhibits this effect. When KaiC is completely unphosphorylated, the KaiB influence releases and the cycle can begin again. Because this cycle repeats daily under endogenous control, the molecular clock in *S. elongatus* is a circadian mechanism. Because the clock's output regulates other molecular systems, this mechanism is causally relevant to the circadian rhythmicity of phenomena in *S. elongatus*.

Although details remain to be worked out, this basic mechanistic account of the KaiC cycle is well-established (see Mackey, Golden, & Ditty, 2011, for a review). Paddock et al. convey

[2] In most organisms, the clock has been understood as a transcription-translation feedback loop, in which for certain specialized genes the translation into mRNA, and the mRNA's transcription into protein, occur on a 24-hour cycle due to their organization as interacting positive and negative feedback loops. This contrasts with the posttranslational clock for *S. elongatus*, which involves changes of state in the protein rather than rhythmicity in the genetic mechanism generating it.

the sequence of states in a mechanism diagram employing "glyphs" (Tversky, 2011)—here, iconic symbols for the parts and arrows for the operations—in a particular spatial layout (Figure 1 here and in Paddock et al.). The glyph for KaiC conveys its structure as a dual hexamer. Solid arrows indicate the phosphorylation and dephosphorylation operations that alter the state of the KaiC component. Spatial contiguity is used to indicate when KaiA and KaiB are interacting with KaiC to bias those operations. The resulting sequence of four states, starting with unphosphorylated KaiC at the far left (dawn) is laid out in a circle to make their cyclic temporal organization immediately apparent.

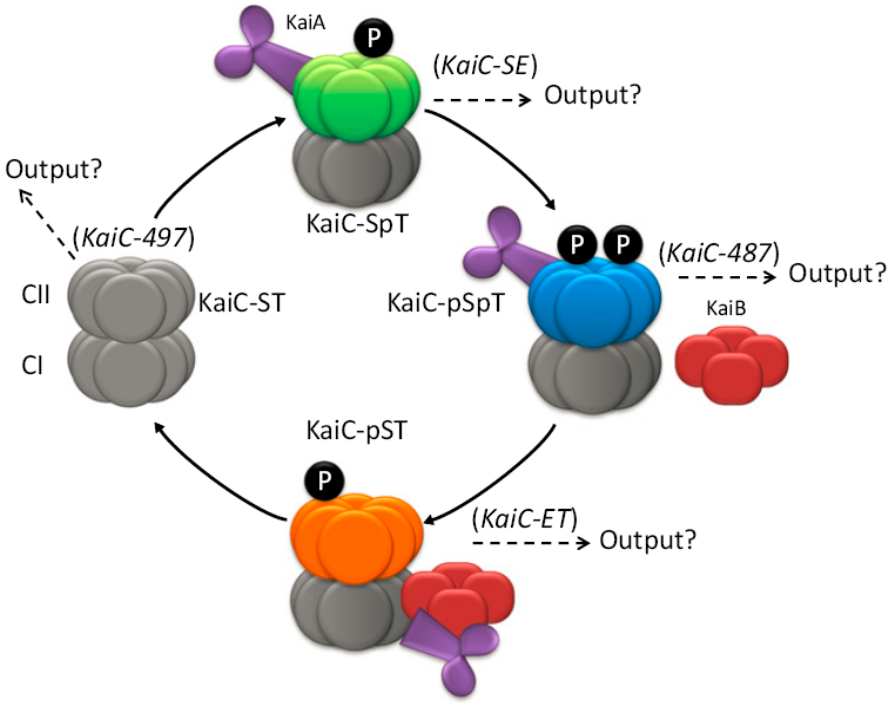

Figure 1. The first figure from Paddock et al. (2013), showing the accepted central clock mechanism for *S. elongatus*. Kai C is the large dual-hexamer molecule. KaiA is the purple molecule and KaiB the red tetramer. Black circles labeled "P" represent phosphate groups, and each phosphorylation state of KaiC is both labeled (inside the circle) and color-coded (on the top hexamer). The first version of this diagram was dated November 17, 2012; the first version that had four "Output?" arrows to indicate uncertainty was dated March 13, 2013. See the text for more detail.

The researchers' inclusion of four dashed arrows terminating in "Output?" draws attention to their focus on the *output* of the central clock mechanism rather than its already-established internal components. If the endogenously generated clock rhythm is to be of any biological value, some output signal(s) must connect the clock causally to other systems, transmitting its rhythmicity to the timing of gene expression in these systems. At the time of the paper's conception, there was a standard mechanistic model for clock output in *S. elongatus* that involved two additional proteins, SasA and RpaA. On this model, as shown in the mechanism diagram in Figure 2, a single output pathway from KaiC to SasA

to RpaA produces a cascade of rhythmic effects —all prominently involving phosphorylation states—culminating in the rhythmic expression of genes across the genome.

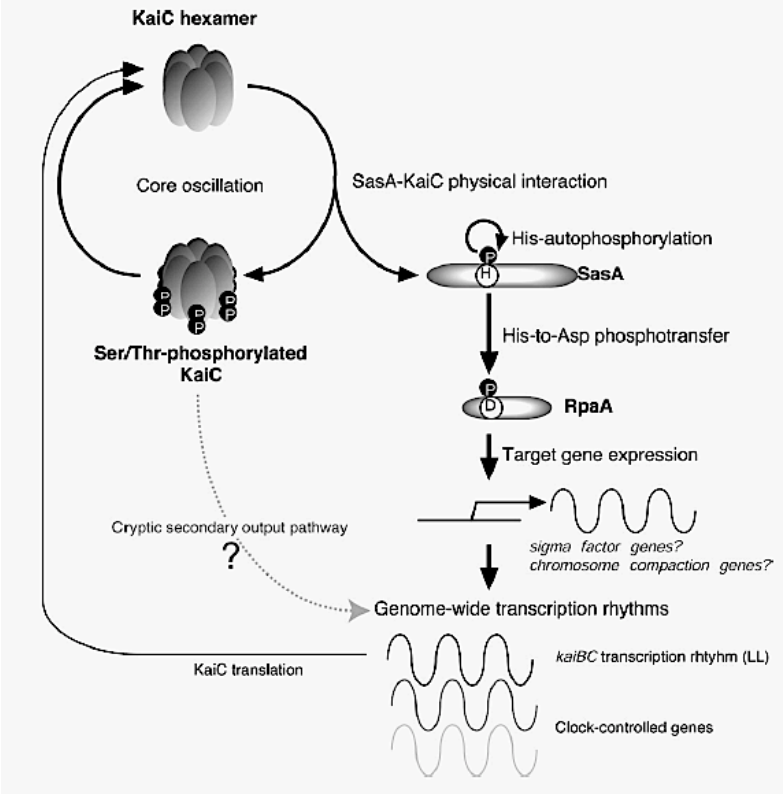

Figure 2. Proposal for the output mechanism from the central clock in *S. elongatus*: from Figure 4 in Takai, Nakajima, Oyama, Kito, Sugita, Sugita, Kondo, & Iwasaki (2006). Note that it does not identify the exact phosphorylation stage of KaiC at which output is sent from the central oscillator through the SasA/RpaA pathway. Note too the "cryptic secondary output pathway," which was included because the researchers had identified residual gene expression rhythms when SasA and RpaA were not present.

The impetus for Paddock et al.'s research was the recent availability of techniques enabling more fine-grained investigation of the output mechanism. As we discuss in some detail below, this work resulted in the discovery of new explanatory relations that led to their reconceptualization of the mechanism. The standard output model had been developed to account for the results of experiments on knockouts lacking SasA or RpaA. SasA knockouts exhibited a severe reduction of rhythmicity in reporter gene expression despite intact rhythmicity in the central KaiC clock. This suggested the explanatory relevance of SasA in output signaling (Iwasaki, Williams, Kitayama, Ishiura, Golden, & Kondo, 2000). Similar results were obtained later with RpaA knockouts, and the ordering of the output operations (from SasA to RpaA) was inferred from SasA's (but not RpaA's) ability to bind DNA.

However, the standard model left several questions unanswered—most notably, which phosphorylation state of KaiC generated the temporal output signal (thus the four "Output?" labels in Figure 1 and the unspecified state of KaiC that physically interacts with

SasA in the output in Figure 2 above). Since KaiC was present in all of its phosphorylation states in the studies establishing the standard model, it was impossible to tell. Paddock et al. hence initiated new experiments that not only yielded an unexpected answer to this question but also pushed them towards proposing an output mechanism quite different than the standard model.

To pursue the question of which KiaC phosphorylation state served as the central clock's output, Paddock et al. employed two techniques, one experimental and one analytic. For their experiments they constructed four variants of KaiC, each of which could stably mimic *in vivo* the behavior of KaiC in one of its four phosphorylation states. Next to each state in Figure 1, in parentheses, is the name given its phosphomimetic ("mimetic" for short). Each of the mimetics could then be introduced individually against background KaiC knockouts to determine its effect, if any, on gene expression. The pattern of results across the four mimetics would provide the key explanatory relation for inferring which KaiC state (or states) is the output source. Moreover, since the effects were not expected to be the same across the entire genome, they measured the activity of two types of gene promoters—"class 1" and "class 2" promoters—each of which activates transcription of a different large set of *S. elongatus* genes. In wildtype bacteria, class 1 promoters tend to be most active at dusk and class 2 promoters at dawn, suggesting that they indeed are regulated by the clock via its output pathway. The researchers selected P_{kaiBC} as an exemplar of class 1 and P_{purF} as an exemplar of class 2. The activity of each promoter was measured by inserting into the DNA region it controlled a bioluminescence reporter gene, *luciferase (luc*). When the promoter activates *luc*, it synthesizes the luciferase protein, which in turn catalyzes a chemical reaction in which light is emitted. The resulting fluctuations in luminescence across the dozens of hours of an experiment were readily tracked and quantified, providing a window on each promoter's activity that was not available from the native genes it regulates.

The researchers' goal was to determine how phosphorylation states of KaiC relate to gene expression governed by the two types of promoters. As we will see in the next two sections, this involved determining the most relevant explanatory relations *via* a novel analytical technique; devising diagrams that best displayed these relations visually; and figuring out implications for how the mechanism was organized. This occurred over several months after all the data had been gathered. In the end, the researchers identified two new explanatory relations, each of which supported a partial reconceptualization of the output mechanism. In both cases, new explanatory relations diagrams and new mechanism diagrams were integral to the research process.

3. Recharacterizing the Output to Determine its Source

The first set of findings zeroed in on determining which of KaiC's four phosphorylation states was the source of the clock's output. The relevant manipulations involved knocking out the KaiC gene and then introducing each of the four mimetics in turn—a way of isolating the effects of each state by proxy. As shown on the left side of Figure 3, the researchers initially graphed bioluminescence (counts per second) across 144 hours for

each condition to be compared. It can first be noted that the data patterns were different, but complementary, for class 1 promoters (panel A) and class 2 promoters (panel B). For both promoters, the baseline bioluminescence pattern is the circadian oscillation found in wildtype bacteria (as indicated by the oscillating pattern of the black square datapoints). The KaiC knockouts (black diamonds), lacking the oscillator mechanism shown in Figure 1, exhibit arrhythmia; their default promoter activity is high with the class 1 reporter and low with the class 2 reporter.

The key question was which of the phosphomimetics, when added to the knockout, would produce an effect. It can be seen that only KaiC-ET (orange squares) yielded data substantially different than the knockout; the other three mimetics, shown here just for KaiC-SE (green circles), produced data similar to the knockout and hence were judged to be ineffectual. The effect of KaiC-ET was substantial: it reduced bioluminescence with the class 1 reporter to the level of the wildtype trough (panel A), and increased bioluminescence with the class 2 reporter to the level of the wildtype peak (panel B). KaiC-ET is the mimetic for KaiC-pST, the state for which the S site, but not the T site, is phosphorylated. The fact that KaiC-ET, and none of the other mimetics, has a quantitative effect on promoter activity is thus a key explanatory relation. It picks out KaiC-pST (bottom of Figure 1) as the state that produces output. This is especially interesting in that previous researchers had suspected that the fully phosphorylated KaiC-pSpT was the output source.

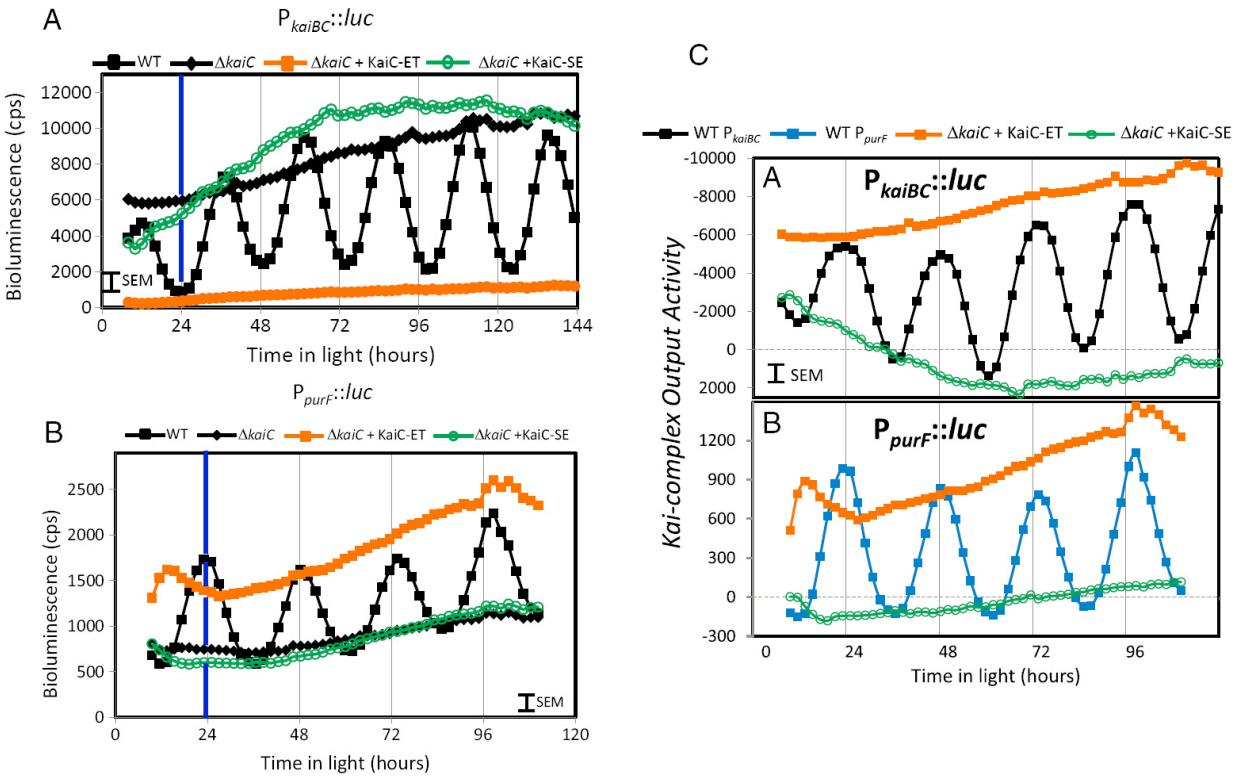

Figure 3: Panel A: Paddock et al.'s Figure 3C. Panel B: Paddock et al's Figure 4. Panel C: Paddock et al.'s Figure 6. All three were first developed on November 17, 2012. See text for details.

The bioluminescence graphs show that there are explanatory relations between KaiC-pST and both types of promoter. The specific quantitative data patterns, showing the contrast between WT, KaiC knockout, and KiaC-ET conditions afford particular inferences in reasoning about the mechanism. In particular, the fact that KaiC-ET bioluminescence levels are the opposite of levels generated by the knockout, and equal to either the peak (class 1) or trough (class 2) of WT levels, suggests not only that KaiC-pST is what is responsible for producing those phases in the WT rhythms, but that cycles in KaiC-pST are likely vital for producing the actual oscillation in WT bacteria.

The effects of KaiC-pST on the two promoters appear dissimilar, producing low values in one case and high in the other. To better understand the role of KaiC in promoter activity writ large, Paddock et al. developed new ways of graphing the results that reveal that these apparent differences mask underlying similarities in KaiC effects. They first employed a new measure of the data, which they referred to as Circadian Activity (CA) or Oscillator Output Activity (OOA) in their early drafts, and as Kai-complex Output Activity (KOA) in the published version. Values on this measure are obtained by subtracting the value of the promoter bioluminescence in KaiC knockouts from the value of the promoter bioluminescence in whatever variety is currently of interest (the KaiC wildtype or one of the phosphomimetics). This new measure specifies the *change* in bioluminescence relative to that of a bacterium that is similar but for the loss of its molecular clock—a sensible measure of the overall output of each oscillator of interest. Notice that calculating KOA does not add anything to the data and is not a new experimental method, but rather an analytical tool for representing certain aspects of the relations discovered. As the authors note, "Although this calculation is merely a transformation of the reporter data, it more clearly reveals the following: (i) the influence of the circadian clock on the two predominant classes of promoters, class 1 and class 2; (ii) the point in the cycle when the magnitude of KOA is the greatest for a given promoter; and (iii) which KaiC variant/phosphostate invokes the greatest magnitude of KOA" (E3853).

Each of these effects are best seen in the graph of KOA in panel C of Figure 3, which plots the KOA measures of the same data shown in panels A and B. Consider first the WT findings. Since the KaiC knockout revealed a default state of *high* activation for the class 1 promoter, it is to be inferred that in WT, KaiC is engaged in *repressing* these promoters' activity. Correspondingly, the value of KOA – the measure of oscillator activity – is highest when activation of class 1 promoters is lowest. As a result, the oscillatory pattern shown in Figure 3 panel A is inverted in the top half of panel C (e.g., the troughs for bioluminescence values reflecting class 1 promoter activity in panel A at 24, 48, 72, . . . hours become the peaks of KOA, as shown in panel C). In contrast, since the KaiC knockout revealed a default state of *low* activity for class 2 promoters, it is to be inferred that KaiC enhances these promoters' activity. As a result, the oscillatory pattern shown in panel B is not much changed in the bottom half of panel C (e.g., the troughs of bioluminescence activity in panel B correspond directly to the troughs of KOA for class 2 promoters).

The new graph affords a certain kind of inference that could only be extracted with some difficulty from the previous panels, namely that while the ultimate effects on the two promoter types are different (resulting in a high value in one case and a low value in the

other), KaiC exerts a quantitatively similar level of influence on each at similar times during the day. It shows that the intact central oscillator is always most active in regulating gene expression at dawn (24 h, 48 h, etc., the times when the WT most deviates from KAO=0), and that this is true for both promoters. Importantly, the contrast between the panels presenting raw data and those presenting KOA shows explicitly that the oscillator's opposite effects on class 1 and class 2 promoters (repressing the former and enhancing the latter, indicated by the fact that the values on the y-axes are negative for class 1 and positive for class 2 promoters) show a similar quantitative and temporal pattern of influence. Finally, it demonstrates that the clock's output is dependent specifically upon the KaiC-pST state (indicated by the fact that the KaiC-pST mimetic, KaiC-ET, closely replicates the peak outputs of the oscillations exhibited in the WT for both class 1 and class 2 promoters) and cannot be recovered by other KaiC mimetics (such as the KaiC-SpT mimetic shown in the green line for contrast). The KOA analysis thus exhibits the importance of data organization for inference, in that a different organization of the data makes different aspects of the relationships available for inference.

Having established that (a) KaiC-pST is the explanatorily relevant phosphorylation state of KaiC for generating the output signal and (b) there are two output effects, excitatory and inhibitory, the researchers returned to their initial mechanism diagram (Figure 1) and undertook several rounds of revision. We discuss two revisions here (and substantial further revisions in Part 4). First, in December 2012 they produced the mechanism diagram in panel A of Figure 4 below. Having resolved the uncertainty indicated in Figure 1 by the four dashed "output?" arrows, the researchers now showed the Kai-pST state as the sole source of output. Moreover, their findings enabled them to replace the dashed arrow with two arrows that distinguish between the output's enhancement of activation for the class 2 promoter (standard arrow) and its inhibitory effect on the class 1 promoter (flat-headed arrow). The mechanism diagram thus captures two inferences about the organization of the output mechanism, which are afforded by the quantitative results displayed in the explanatory relations diagrams in Figure 3. From the finding with KAO that the KaiC-pST mimetic (KaiC-ET, in panel C) exerts substantial quantitative effects on the two types of promoters, it was inferred that KaiC-pST is the origin of the output pathway. From the finding that adding the KaiC-pST mimetic to the KaiC knockout results in trough expression for class 1 (panel A), but peak expression for class 2 (panel B), it was inferred that KaiC-pST inhibits the one and activates the other. Once these inferences were captured in a revised mechanism diagram, it could be seen where the explanatory relations linked into it. In the second round of revision a month later (Figure 4B), the researchers foregrounded the explanatory relations by omitting the glyphs for the three ineffective states, leaving just a circle to represent the central KaiC clock, but adding labeled flows of input and output.

Even after developing the mechanism diagram in Figure 4B, the researchers spent months developing further variants and finally a major reconstrual of the mechanism. Why? It turns out that there was more to explain. In the next section, we examine this last stage in their co-development of explanatory relations and mechanism diagrams.

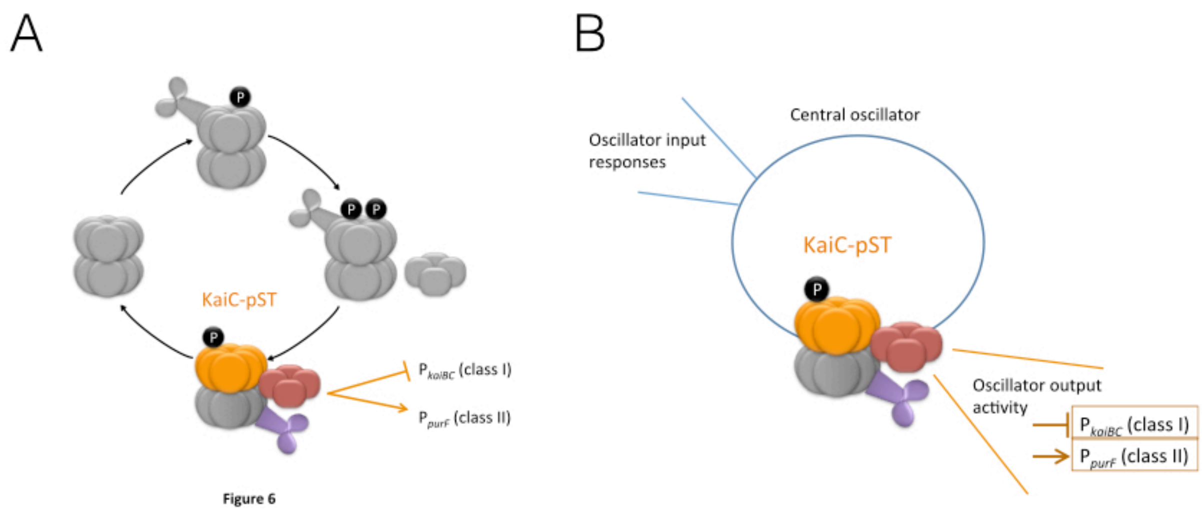

Figure 4. Panel A: an early draft of what was published as Figure 7, dated December 4, 2012. Panel B: a pared down version of the figure dated January 11, 2013. See text for details.

4. Rethinking the Role of RpaA

In early drafts of the Paddock et al. manuscript, the results discussed in the previous section were presented as the primary conclusions of the paper. The mechanism diagrams in Figure 4, for example, focused on the central clock and included only a minimalist representation of the output mechanism. In particular, they did not indicate a role for RpaA, which figured importantly in the standard model shown in Figure 2. Might RpaA nonetheless be part of the pathway through which KaiC-pST regulates gene expression? That would be the simplest way to integrate their findings about KOA with the standard model. However, Paddock et al.'s initial data had included results from knockouts of RpaA and measurements of its phosphorylation state in the different KaiC conditions that could be mined to address this question. They puzzled this out, especially in March-April 2013, during which they generated new explanatory relations diagrams and a series of mechanism diagrams that together resulted in a radical reconceptualization of the output mechanism.

Already in their first manuscript of October 2012, Paddock et al. had produced an earlier version of Figure 5C below. It used immunoblots to show the relative amounts of phosphorylated vs. unphosphorylated RpaA in several strains, including WT, KaiC knockout, and two of the mimetics (KaiC-SE and KaiC-ET). If the output pathway involved transmission of phosphorylation from KaiC-pST to SasA to RpaA, then the ratio of phosphorylated to unphosphorylated RpaA should have been higher in the KaiC-pST mimetic (KaiC-ET), than in the other strains. The finding that the ratio of phosphorylated to unphosphorylated RpaA was basically the same in all four conditions shows, however, that

this predicted explanatory relation does *not* hold and therefore something about that proposed pathway was incorrect.

When Paddock et al. finally pursued this lead, they took advantage of the fact that they had included an RpaA knockout in the experiments measuring bioluminescence as a proxy for gene expression. This enabled them to generate explanatory relations diagrams comparing the RpaA knockout to the KaiC knockout, from which they could make inferences about the output mechanism. Eventually published as Panels A and B of their Figure 5, these diagrams revealed that for both classes of promoters, an RpaA knockout (red squares) had effects on bioluminescence that were the reverse of those for the KaiC knockouts (black diamonds). In the October draft this finding was mentioned only as "data not shown." In a November draft the diagrams in Panels A and B, along with Panel C, were designated to appear in an online supplement, and in a December draft they were elevated to the main article. Finally, in March 2013 the researchers added panel D. This shows that, with RpaA deleted, KaiC still affects the class 2 promoter albeit with a loss of oscillation: with KaiC present bioluminescence is high (red squares) and with KaiC deleted bioluminescence is low (blue squares). This explanatory relation, along with others, pointed to an output mechanism comprising two pathways, one of which was independent of RpaA.

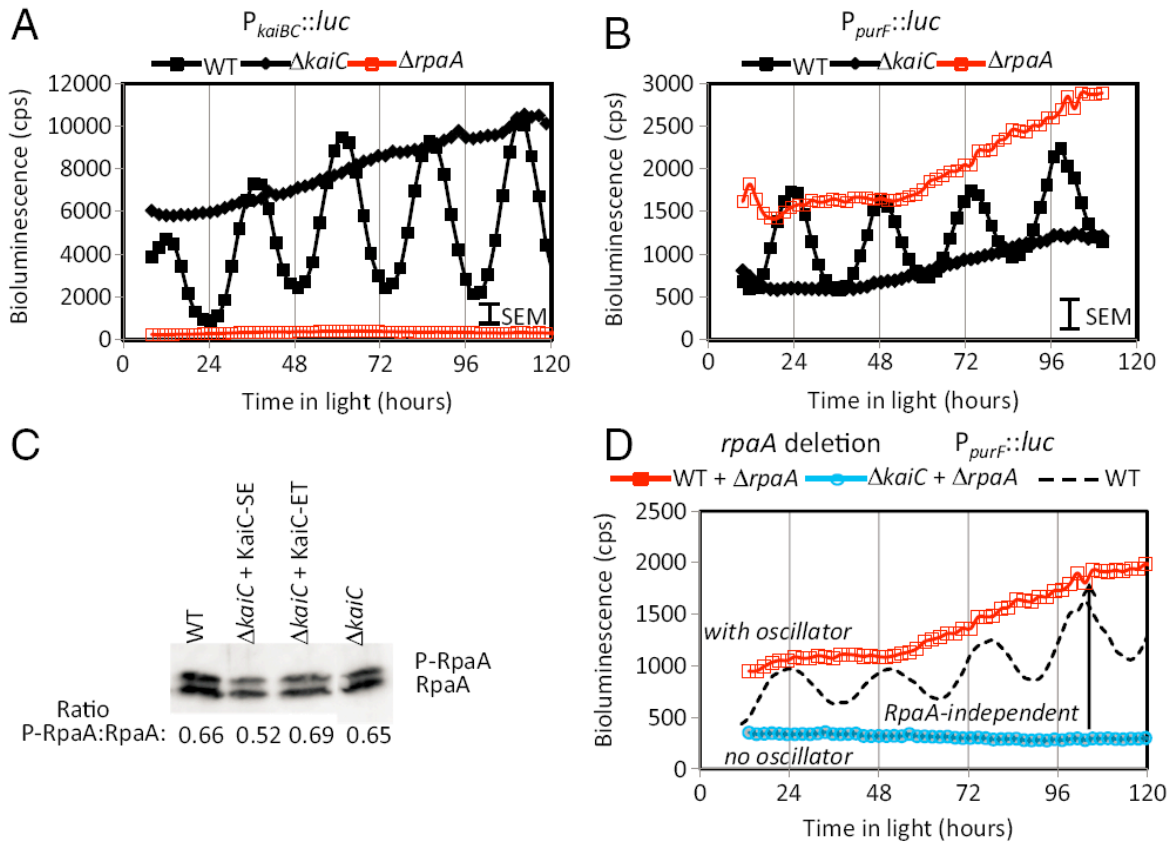

Figure 5: Paddock et al.'s Figure 5. Panels A and B: First developed on December 4, 2012. Panel C: First developed on October 21, 2012. Panel D: First developed on March 7, 2013. See text for details

Since the researchers did not generate new data during this time, the above process is best described as one in which by March 2013 the researchers' focus had shifted to seeking new explanatory relations in the original data set from which to make inferences about the output mechanism. The researchers repeatedly re-grouped and visualized the existing data to reveal important explanatory relations—e.g., the RpaA-independent effect of KaiC-pST on class 2 promoter activity. After the revised version of Figure 5 was incorporated in the main paper draft in March 2013, including the new explanatory relation in panel D, the proposed scope of the paper expanded to incorporate findings on the relative contributions of KaiC-pST and RpaA and, pursuing the implications, a reconceptualization of the output mechanism.

Starting in March 2013, the researchers generated a number of draft mechanism diagrams; Figure 6A is representative. Their goal was to achieve an overall understanding of the mechanism by incorporating what was known about the RpaA pathway from the standard model—namely its necessity for output, and the importance of its phosphorylation states—along with what they themselves had learned about the RpaA-independent pathway. The researchers proposed that, in its phosphorylated state, RpaA represses the effects of the proposed pathway from KaiC-pST. Notably, this mechanism diagram explicitly includes a reference to oscillator output activity (their earlier term for KOA) and thereby shows how the explanatory relations discussed in section 3 and the posited mechanism organization are related. The dotted line extending from the non KaiC-pST phosphorylation states is meant to indicate that in the absence of the KaiC-pST state, OOA is at its trough levels. The solid arrow from the KaiC-pST state to the peaks of the output activity is intended to indicate that the KaiC-pST state produces maximum OOA.

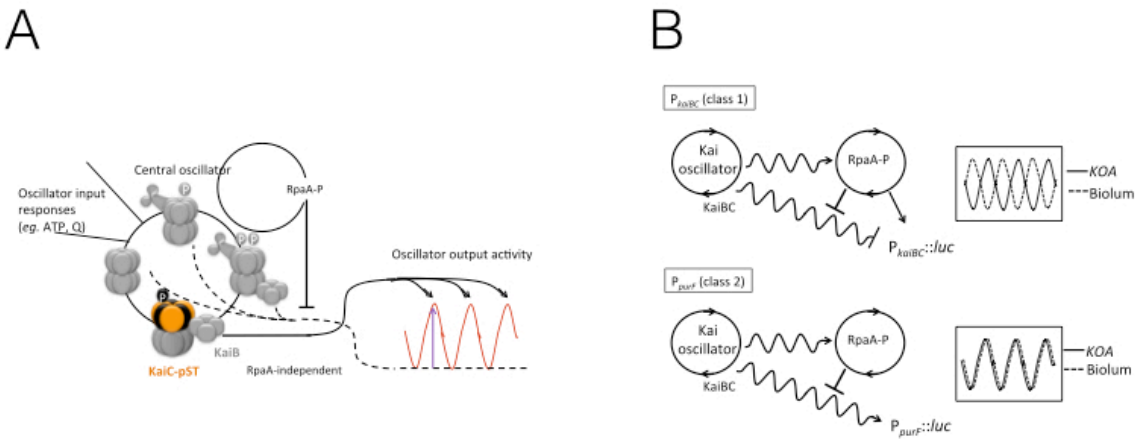

Figure 6. Panel A: one of several intermediate versions of the left half of Paddock et al.'s Figure 7 that appeared in versions of the manuscript from early March until mid-April 2013. Panel B: a pared down version that appeared on April 11, 2013. See text for details.

This figure shows phosphorylated RpaA *inhibiting* the proposed output pathway from KaiC-pST, which is inferred from the explanatory relations shown in figure 5 (panels A and B)—that RpaA and KaiC have opposite effects on promoter activation. Not shown are the differential effects of that pathway on class 1 vs. class 2 promoters. Other versions developed during the same period (not shown here) tried a strategy like that of Figure 4 to include those effects, but it was clearly a challenge to display both pathways, their interaction, and differential effects in a single mechanistic diagram. This, plus the visual complexity of using icons to convey the different phosphorylation states of KaiC in the same figure, may have been what led to a rather major revision on April 11, as shown in Figure 6B.

This new diagram of the output mechanism abandons prior attempts to portray the new conceptualization by extending the iconic design that worked well in the simpler Figure 1. It represents only one state in the KaiC cycle explicitly, now called KaiBC-pST in recognition of the role of KaiB, and clearly shows it is responsible for the output to the newly proposed pathway. The KaiC state(s) responsible for the output to the original pathway through RpaA is shown as still indeterminate, but the fact that its outputs are oscillatory is now indicated iconically. Panels B differentiates the cases of the class 1 (top) and class 2 (bottom) promoters. The pathway from KaiC-pST to the promoter is shown as inhibitory in the case of class 1 promoters. When the central oscillator is in the KaiC-pST phosphorylation state and RpaA is not phosphorylated, KaiC-pST represses the promoter, as is revealed in the explanatory relation of KOA being out of phase with bioluminescence. When RpaA is phosphorylated (RpaA-P) it both represses the output from KaiC-pST and enhances KOA. Thus, the interaction of the central oscillator and the oscillation between RpaA and RpaA-P generate the oscillation in KOA in class 1 promoters. In the case of class 2 promoters, the effect of Kai-pST is excitatory, reflecting the fact that, as shown in Figure 5b, both bioluminescence and KOA levels were locked at high levels in the RpaA knockout. When RpaA is present and phosphorylated it represses this output, generating the observed oscillation. Note the radical difference between this organization of the mechanism and the serial output pathway shown in Figure 2.[3]

Finally, in May 2013 the authors arrived at a version that appeared in the published version with only minor changes (Figure 7). Unlike in the previous versions, KOA is not shown in the same panels as the mechanism, but instead in panel C, where values from Figure 3C above are relativized to the percentage of maximal expression of the WT. Thus, the green line shows the low output of the KaiC-SpT phosphorylation state, the gold line the high output of the KaiC-pST state, and the blue line the WT values. For comparison, the KOA for the RpaA mutant is shown in a separate panel at the top; notice that KOA is much greater in the RpaA knockout, highlighting its role as a repressor. As indicated by the

[3] The researchers do include a direct arrow from RpaA to the class 1 promoter. As they note, the greater activation of class 1 promoters due to RpaA is consistent with the standard model, which in this case needs to be extended by introducing the independent KaiC-pST pathway and the inhibition from RpaA. The class 2 results, they also note, are more straightforwardly incompatible with the standard model and hence there is no arrow to the class 2 promoter.

equation below the graph, the authors also developed a computational model of the mechanism presented in panels A and B; the red line in the graph shows the results of a simulation using the model. We cannot discuss the simulation further here other than to note that one important role for such simulations is to show that a proposed mechanism could actually generate the dynamic behavior attributed to it.

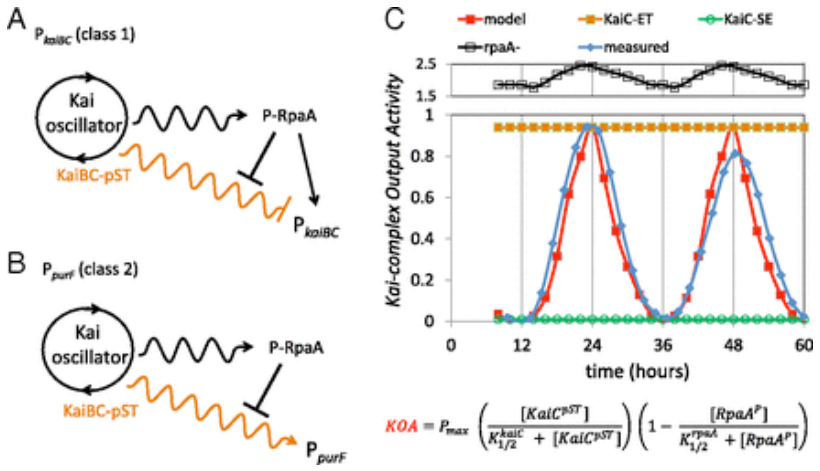

Figure 7. The published version of Figure 7 in Paddock et al. (2013).

While in the final version KOA is shown in a separate panel from the mechanism, its inclusion next to the mechanism diagram in Figure 7 exemplifies an important fact about representational practice: mechanism diagrams are not stand-alone explanations. Instead, they implicitly or explicitly cooperate with diagrams of explanatory relations in generating an understanding of the mechanism and how it produces the phenomenon of interest. Explanatory relations diagrams should not be viewed as just a stepping-stone to developing mechanism diagrams. In fact, full understanding requires relating the mechanism diagram to the relevant explanatory relations diagrams such as those shown in Figure 3. The summative figure of the paper thus includes both kinds of representations in an explanatory role, evidencing the cooperative relationship between these two kinds of diagrams.

5. Conclusion

Scientific reasoning is both important and complex. As such, it can be studied using a variety of methods and conceptual frameworks. Our methodology in this paper was to examine representational tools constructed by scientists, specifically diagrams, and their development in the course of a research endeavor from various unpublished versions to the final products. This enabled us to gain insight into the reasoning that leads to scientific conclusions. The two kinds of diagrams provide different resources for explanation. Mechanism diagrams display parts and operations, laying them out in space so as to indicate spatial, temporal, and/or causal organization that can provide a basic (nonquantitative) explanation of a phenomenon. Explanatory relations diagrams extract comparisons from data that may show the need to add, delete or change certain mechanism parts or operations, while adding quantitative detail linked to these components that is distinctively, nonredundantly explanatory. Those linkages indicate where each particular

explanatory relation fits into the overall organization. Moreover, the ordered corpus of diagrams of each type can be regarded as an external record that, in tracking changes in the explanations, parallels the scientist's reasoning.

In this case study, both explanatory relations diagrams and mechanism diagrams were repeatedly modified, as the scientists reasoned back and forth between diagrams and between diagrams and their manuscript drafts. Inferences from particular explanatory relations to changes in the mechanistic account were especially salient. In principle, this kind of reasoning could be brought into a cognitive science laboratory and studied with various controls in place and large numbers of participants (see Stieff, Hegarty, & Deslongchamps, 2011, for one such study). It could as well be pursued by philosophers of science in further case studies or ethnographic studies that target interestingly different scientific endeavors. Hence, in exploring the distinct roles of different types of diagrams and how they were coordinated in the course of reasoning towards an explanation, we have provided a target and resource for further analysis in both philosophy and cognitive science.

Acknowledgment

We gratefully acknowledge the support of U.S. National Science Foundation grant 1127640. We also thank Drs. Mark Paddock and Susan Golden for their generous cooperation with this project, including access to their draft documents.

References

Alač, M. (2011). *Handling digital brains: A laboratory study of multimodal semiotic interaction in the age of computers*. Cambridge, MA: MIT Press.

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. A. (2012). Thinking dynamically about biological mechanisms: Networks of coupled oscillators. *Foundations of Science, 18*, 707-723.

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.

Burnston, D. (forthcoming). Data graphs, explanatory relations, and mechanistic explanation.

Dunlap, J. C., Loros, J. J., & DeCoursey, P. J. (Eds.) (2004). *Chronobiology: Biological timekeeping*. Sunderland, MA: Sinauer Associates.

Hegarty, M. (2011). The cognitive science of visual-spatial displays: Implications for design. *Topics in Cognitive Science, 3*, 446-474.

Iwasaki, H., Williams, S. M., Kitayama, Y., Ishiura, M., Golden, S. S., & Kondo, T. (2000). A kaiC-interacting sensory histidine kinase, SasA, necessary to sustain robust circadian oscillation in cyanobacteria. *Cell, 101*, 223-233.

Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science, 67*, 1-25.

Mackey, S. R., Golden, S. S., & Ditty, J. L. (2011). The itty-bitty time machine: Genetics of the cyanobacterial circadian clock. *Advances in genetics, 74*, 13-53.

MacLeod, M., & Nersessian, N. J. (2013). Coupling simulation and experiment: The bimodal strategy in integrative systems biology. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences, 44*, 572-584.

Osbeck, L. M., Nersessian, N., Malone, K. R., & Newstetter, W. C. (2010). *Science as psychology: Sense-making and identity in science practice*. Cambridge: Cambridge University Press.

Paddock, M. L., Boyd, J. S., Adin, D. M., & Golden, S. S. (2013). Active output state of the Synechococcus Kai circadian oscillator. *Proceedings of the National Academy of Sciences, 110*, E3849-E3857.

Perini, L. (2005). Explanation in two dimensions: Diagrams and biological explanation. *Biology and Philosophy, 20*, 257-269.

Sheredos, B. (2013). Scientific diagrams as traces of group-dependent cognition: A brief cognitive-historical analysis. *Proceedings of the 35th Annual Meeting of the Cognitive Science Society* (pp. 2396-2401). Austin, TX: Cognitive Science Society.

Sheredos, B., Burnston, D., Abrahamsen, A., & Bechtel, W. (2013). Why do biologists use so many diagrams? *Philosophy of Science, 80*, 931-944.

Stieff, M., Hegarty, M., & Deslongchamps, G. (2011). Identifying representational competence with multi-representational displays. *Cognition and Instruction, 29*, 123-145.

Tversky, B. (2011). Visualizing thought. *Topics in Cognitive Science, 3*, 499-535.

Woody, A. I. (2000). Putting quantum mechanics to work in chemistry: The power of diagrammatic representation. *Philosophy of Science, 67*, S612-S627.