

From Reduction Back to Higher Levels

William Bechtel (bechtel@mechanism.ucsd.edu)

Department of Philosophy-0119, UCSD
La Jolla, CA 92093-0119 USA

Adele Abrahamsen (aabrahamsen@ucsd.edu)

Center for Research in Language-0526, UCSD
La Jolla, CA 92093-0526 USA

Abstract

In the context of mechanistic explanation, reductionistic research pursues a decomposition of complex systems into their component parts and operations. Using research on circadian rhythms and memory consolidation as exemplars, we consider the gains to be made by finding genes and proteins that figure in mechanisms underlying behavioral phenomena. However, we also show that such research is insufficient to explain the initial phenomenon. Accordingly, researchers have increasingly recognized the need to consider higher-level organization and integration with other systems. This illustrates a common need to complement reductionistic inquiry with investigations at higher levels and identifies a trajectory whereby cognitive science can embrace molecular neuroscience without surrendering its own contributions.

Keywords: reduction, mechanistic explanation, memory consolidation, circadian rhythms.

Introduction

The rise of cognitive neuroscience offers both opportunities and challenges to cognitive scientists. In addition to neuroimaging and other new tools for linking cognitive processes to brain regions, it opens a potential conduit to the thriving fields of cell and molecular neuroscience. What should cognitive scientists make of this? To many it brings the threatening prospect of accounts of human behavior in terms of genetic and biochemical processes, leaving little room for cognitive scientists' theoretical and computational models.

This is, in fact, how reductionistic research is often portrayed by both its advocates and critics (see papers in Schouten & Looren de Jong, 2007). The goal of reduction is seen as completely explaining the phenomenon of interest at the lowest possible level (e.g., in terms of genes and biochemistry), thereby supplanting and rendering superfluous the kinds of accounts typically offered by cognitive scientists or even those of systems neuroscientists (Bickle, 2003). That is, if one can account for and predict all that happens in terms of the lowest-level parts and operations, there is no need for any additional account at a higher level. A researcher invoking psychological processes, for example, is trying to explain what has already been explained. The psychological narrative is at best epiphenomenal (that is, psychological processes result from the lower-level processes and have no causal efficacy of their own).

We will argue that this seriously misrepresents reductionist research, which even when most successful does not provide a complete account of the phenomenon of original interest. It uncovers crucial components, but these must be fit

into a more comprehensive account that considers processes at the higher levels that were initially left behind in the reductionistic quest. To trace both the downwards and upwards trajectories in specific cases, we examine two reductionistic research programs targeting behavioral phenomena. In both cases the reductionistic pursuit has been highly successful—an enormous amount has been learned about the genes and biochemical processes involved. But this success has been very local, and itself points to the need to integrate what has been learned into higher-level accounts.

One of these cases involves research on memory consolidation, a phenomenon identified and studied by psychologists beginning in the late 19th century, but investigated primarily at the molecular level since the 1970s. The other case, circadian rhythms, also began with behavioral studies (by evolutionary and behavioral biologists), with pursuit of the molecular level added more recently. Molecular biologists working in these domains make use of others' ongoing investigations at higher levels. This suggests that investigation of the neural processes underlying more prototypically cognitive domains likewise will require complementation by higher-level inquiries such as those pursued by cognitive scientists. On this view, reductionism is not a threat to replace cognitive accounts; instead, it leads to new information that can enrich and improve those accounts.

The contrasting view of reduction, in which lower-level accounts supplant higher-level ones, is anchored in a 20th century philosophy of science that emphasizes laws as the explanatory engine. From this perspective, if laws existed that completely characterized how lower-level entities behaved in all contexts, it is hard to see what a higher-level account could add (Kim, 1998). But it is unclear what these laws would be like, since current laws in physics only characterize the behavior of physical objects in highly idealized contexts in which they are isolated from other factors that usually impinge on them (Cartwright, 1999). The applicability of law-based accounts to the biological and cognitive sciences is dubious as well, since explanations in these sciences seldom invoke laws. Rather, they most frequently take the form of identifying the mechanism responsible for a given phenomenon. Philosophers focused on these sciences have recently articulated a new mechanistic philosophy of science that is especially appropriate to these sciences (Bechtel, 2008; Bechtel & Richardson, 1993; Machamer, Darden, & Craver, 2000; Thagard, 2006).

From the mechanistic perspective, to explain a phenomenon is to explicate the mechanism responsible for it. In each

research area considered to date, this has been achieved by decomposing the mechanism into its parts and operations and showing how these are organized such that the whole mechanism, in appropriate contexts, produces the phenomenon of interest. Whereas the law-based perspective on explanation regards laws at higher levels as derived from those at lower levels, thereby denying them any autonomous function, the mechanistic account attends to the manner in which lower-level processes are organized into higher-level systems which then engage other systems. Scientific inquiry at these higher levels contributes equally with inquiry at lower levels in providing a full mechanistic explanation.

One of the intriguing features of mechanistic research is that even while a reductionistic decomposition is still being pursued, the need for a return to high-level inquiries often becomes obvious, prompting researchers to turn some of their attention to the organization of the components and how the whole mechanism functions within a wider context.

In the following sections we characterize the typical trajectory in mechanistic research, in which scientists begin by delineating and exploring a phenomenon of interest, then seek to uncover the responsible mechanism by decomposing it into parts and operations (sometimes iteratively), specify which parts perform which operations, and then progressively ascend back up to higher levels (again, sometimes iteratively). We will illustrate both the path down to the molecular parts of a mechanism and the path back up to the system first with recent research on circadian rhythms (in which the behavioral phenomena are relatively circumscribed) and then with research on memory consolidation (in which the behavioral phenomena are more complex, but also more familiar to cognitive scientists).

Delineation and Exploration of the Phenomena

Inquiry typically begins with the discovery and systematic characterization of one or more related phenomena of interest, often followed by decades of exploration in which the account becomes both broader (additional tasks, species, organs, and related phenomena) and more refined (additional detail, corrections, and incorporation of variations).

Circadian rhythms. Observations that living systems perform different activities at different periods of day have been made for centuries, but in the 20th century Pittendrigh (1960) demonstrated that this behavior was often controlled by processes internal to the organism. The phenomenon of circadian rhythms was thus delineated in terms of activity with an approximately 24-hour cycle due to an endogenous “clock.” Researchers have further explored this phenomenon by uncovering the diversity of life forms in which various activities exhibit circadian cycles, notably cyanobacteria, various fungi (*Neurospora*) and plants (*Arabidopsis*), and numerous insects, birds, and mammals. Other researchers have focused on the range of activities exhibiting circadian cycles, identifying not only a number of observable behaviors but also physiological and cognitive functions.

Memory consolidation. Delineation of the phenomenon of memory consolidation is rooted in Ebbinghaus’ pioneering

studies in the 1880s establishing time courses for learning and forgetting nonsense syllables. Later researchers adapted his methods for further exploration of these phenomena and discovery of related ones. Müller and Pilzecker, using their new technique of paired associate learning, in 1900 found evidence for rehearsal and posited that it helped consolidate the linkage between items within each pair. McDougall and Burnham applied Müller and Pilzecker’s work towards explaining what Burnham (1903) called *retroactive amnesia*, the loss of memory for the period immediately preceding shock or injury. This was taken as evidence for a process of memory consolidation that extended across time.

The Path Down to Parts and Operations

The quest to understand the mechanism responsible for a given phenomenon requires decomposing the responsible system. Decomposition typically comes in two flavors: structural (identifying component parts) and functional (differentiating component operations). Such research requires development of techniques that can segregate the parts and reveal the operations. Moreover, as we shall illustrate for both circadian rhythms and memory consolidation, decomposition is often pursued iteratively as investigators tease apart the brain areas, neurons, and ultimately genes and proteins that figure in the phenomenon to be explained.

Circadian rhythms. Circadian researchers were quick to embark on their downwards path. Within a decade of the recognition in 1960 that organisms keep time endogenously, the responsible mechanisms had begun to be characterized at the intercellular (brain area) level and also the intracellular (genetic and molecular) level. Specifically, lesion studies in the early 1970s pinpointed the locus of the primary clock—the *central oscillator*—in mammals as the suprachiasmatic nucleus (SCN), a bilateral cluster of neurons in the anterior hypothalamus just above the optic chiasm.

Identification of components at the molecular level was first achieved in *Drosophila* (fruit flies), not mammals. By using mutagens to produce flies with altered and absent circadian rhythms, Konopka and Benzer (1971) identified the first “clock gene,” *period* (*per*). It was easily linked to basic component operations: like other genes, *per* functions as a template for the production of proteins, with *per* mRNA as mediator in a complex process of protein synthesis that churns out molecules of PER. Once *per* was cloned in the 1980s, it could be shown that concentrations of both *per* mRNA and PER rise and fall on an approximately 24-hour cycle. By the 1990s a mammalian homolog—soon recognized to involve three genes, not one—was identified in mice and humans: *Per1*, *Per2*, and *Per3*. (Note: it is customary to write gene names in italics, with the first letter in lowercase for *Drosophila* and in upper case for mammals, and corresponding protein names in uppercase Roman. Both typically are abbreviated to three letters.)

This was not sufficient to explain how PER cycles. By the mid-1990s the search for more of the clockworks led to the discovery of a host of other clock genes and proteins (the most important shown in Figure 2 below).

Memory consolidation. The downwards path towards a mechanistic, reductive explanation of the phenomenon of memory consolidation was similar in many respects. Neurosurgeon William Scoville focused attention on the role of the hippocampus when he resected it to control epilepsy in the patient HM in the mid-1960s, with the result that HM was no longer able to encode new episodic memories. Notably, HM's pattern of retrograde amnesia for events in the months prior to surgery suggested that a temporally extended memory consolidation process played a crucial role.

Understanding of component mechanisms developed later (and is still ongoing). Donald Hebb is credited with the influential idea that encoding memories involves changes at the neuronal level such that (as later expressed) "neurons that fire together wire together." This would require the generation of new proteins, and Agranoff, Davis, and Brink (1966) provided evidence by showing that administering a protein synthesis inhibitor eliminated learning. Lømo and Bliss (1973) showed enduring changes in the responsiveness of cells in the dentate gyrus area of the hippocampus in rabbits following stimulation sufficient to get the cell to fire, a phenomenon at the level of the neuron that came to be known as *long-term potentiation* (LTP). Though their research had not been directed at studying memory (Craver, 2003), they interpreted the results as suggesting an encoding mechanism.

Researchers quickly took a further downward step to seek out the intracellular processes underlying LTP, which many assumed involved changes in the post-synaptic cell when it produced an action potential following stimulation. They found a host of biochemical constituents which underwent such changes. First, the neurotransmitter glutamate was shown to excite neurons exhibiting LTP. Next, there are two types of receptors in the cell membrane to which glutamate can bind, AMPA and NMDA. AMPA receptors gate the flow of sodium ions (Na^+) into the cell and subsequent outflow of potassium ions (K^+), while NMDA receptors admit calcium ions (Ca^{++}) when the cell produces an action potential. Within the cell are numerous enzymes (e.g., calmodulin) and kinases, including calcium/calmodulin-dependent protein kinase II (CaMKII), protein kinase C (PKC), protein kinase A (PKA) and mitogen-activated protein kinase (MAPK). Researchers also identified specific genes and the molecules that activate their transcription, such as cAMP response element binding protein (CREB).

Upward Path 1: Organization

For both circadian rhythms and memory consolidation, the quest to identify parts and operations that figure in the generation of a behavioral phenomenon led researchers first to brain areas and then to the molecular level of genes and proteins. In both cases, they enjoyed great success in identifying a host of components within one of the subcomponents of one of the major components of the overall system. In the process, the research seems to have left behind the behavioral level at which the phenomena were initially identified. But in both cases the downward path has been com-

plemented by an upward path in which researchers focus on how the components they differentiated are organized into larger systems.

One of the first steps is to identify the complex temporal organization, often involving feedback loops, that links the lowest-level operations into well-orchestrated networks. Spatial organization of component parts is often relevant as well, and a full account will clarify how these are linked. Aspects of organization often are anticipated and sometimes even guide the search for the lowest-level components, but characterizing organization is conceptually distinct from identifying parts and operations. Until they have an understanding of the spatial and temporal organization, researchers find themselves in the same situation as someone who has taken apart a wristwatch, has all the parts laid out on the table, and has no idea how to put it back together. Here we provide a glimpse of how the upwards path was pursued in each of the two fields.

Circadian rhythms. Beyond identifying many, but probably not all of the clock components and reactions they engaged in, researchers sought to understand how these could maintain a reliable 24-hour cycle. The early work focused just on *per* and its protein PER. In *Drosophila*, it was found that concentrations of PER increase in the cytoplasm several hours after the increase in *per* mRNA, and that PER's subsequent transport back into the nucleus occurs just before *per* mRNA levels begin to decline. Harden, Hall, and Rosbash (1990) proposed that PER, once in the nucleus, might somehow inhibit *per* transcription and that this loop between *per* and its protein constituted the core clock mechanism, as illustrated in Figure 1.




Fig 1: Feedback loop: PER inhibits own synthesis

This key feedback loop provided a core to which other component parts and operations of the overall mechanism could be attached; Figure 2 shows the most important. In particular, CKI ϵ phosphorylates the newly formed PER, marking it for degradation, until sufficient CRY is created to form a particular kind of compound with PER (the PER:CRY *dimer*). This then protects PER from immediate degradation, enabling its transport into the nucleus. There it acts upon the CLK:BMAL1 dimer so as to stop it from activating the transcription of *Per* and *Cry*. When CLK:BLAL1 is not acted on by PER:CRY, it binds to the E-box promoter, not only on the *Per* and *Cry* genes but also the *Rev-Erba* gene. REV-ERB α participates in a second feedback loop by inhibiting the transcript of *Bmal1*, thereby reducing the availability of the CLK:BMAL1 dimer that is available to activate *Per* and *Cry* transcription. All these parts and

operations form an integrated system, and it is this system, not the constituents in isolation, that maintains the cycle.




Figure 2: A simplified diagram of the mammalian circadian oscillator

Memory consolidation. A similar process, whereby researchers identified networks of activity involving the lowest-level components, can be seen in research on memory consolidation. The various components involved in LTP—neurotransmitters, receptor molecules, ions, enzymes, etc.—are organized in three different loops, each of which generates increases in the influx of Na^+ into the cell via the AMPA receptors, thereby resulting in a greater depolarization of the cell when glutamate is again released (Figure 3). The most immediate loop involves the Ca^{++} admitted once glutamate binds with the NMDA receptor binding to CaM. This complex in turn binds to its specific kinase CaMKII, which then phosphorylates an AMPA receptor. The result is a change in the configuration of the AMPA receptor that allows a greater influx of Na^+ and, consequently, greater depolarization. In a second loop, CaMKII figures in moving into the membrane additional AMPA receptors that were already created but being held in reserve. These additional functioning AMPA receptors increase the capacity to admit Na^+ and so promote more effect depolarization of the cell. A third loop involves the synthesis of new dendritic spines with their own AMPA receptors. This process utilizes another of the kinases, MARK, which phosphorylates gene transcription factors such as CREB. Few details are known about this process, such as where the resulting mRNAs are utilized to synthesize proteins.




Figure 3: Simplified diagram of three LTP loops

Upward Path 2: Higher-Level Systems

We have seen that researchers focusing on either circadian rhythms or memory consolidation found that the operations carried out within cells by the lowest-level parts of interest (genes and proteins) are organized into complex networks of activity. In the process they begin to appreciate, if only schematically, how such organized systems of components give rise to the phenomenon of interest. But researchers often find that two-level explanations of this kind are still incomplete. The network of genes and proteins itself interacts as a whole with other systems, and understanding these interactions is often a catalyst to moving to yet higher levels of organization, eventually reaching the whole behaving system from which the inquiry was launched.

Circadian rhythms. Research on circadian rhythms has clearly ascended from the initial focus on intracellular mechanisms to higher-level ones, both within the SCN and between the SCN and other organs of the body. We consider first intercellular organization within the SCN. In each hemisphere there are approximately 10,000 interconnected SCN neurons. Do the interconnections contribute to the timekeeping ability of the SCN? By culturing dispersed SCN cells on microelectrode arrays, Welsh, Logothetis, Meister, and Reppert (1995) demonstrated that the rhythms maintained by individual SCN cells, even when still interconnected, vary substantially (SD 1.2 hours). Further research has indicated that each cell is precise in its own distinctive cycle, and yet all SCN cells are entrained daily by light signals (i.e., they are reset and therefore momentarily consistent with each other). In between entrainment events, and in contexts in which entrainment is blocked, the overall output from the SCN more closely approximates 24 hours than does the output of individual cells. That is, an (approximately) 24-hour clock is an emergent property of the interconnected SCN network.

Recent research has pointed to a relationship between the SCN and the rest of the organism that is at least as complex. The initial guiding assumption involved a simple sequential organization: the SCN's input pathway received entrainment signals from photoreceptors (and possibly temperature or other receptors), and the SCN in turn sent cycle-regulating signals along its output pathways to other systems. Hence, the circadian oscillations observed in locomotive, digestive, and other physiological systems were credited to a centralized oscillator in the SCN. When evidence later emerged for the existence of peripheral oscillators in these systems, they were characterized as *slave* oscillators dependent on the SCN's master clock to maintain their rhythmicity. But this simple feedforward picture (the left-to-right sequence in Figure 4) now appears to be false; there is evidence of feedback both from the central oscillator to the entrainment process and from the peripheral systems to the central oscillator (curved arrows in Figure 4).




Figure 4. Feedforward view of the circadian system, modified by adding feedback loops (curved arrows).

Consider first the effects on entrainment. In mammals the entrainment signal involves melanopsin, a photoreceptor in retinal ganglion cells, but there is evidence that the behavior of melanopsin is affected not only by light but also by other cells—including peripheral circadian oscillators located in the retina (Rollag, Berson, & Provencio, 2003). Comparable evidence of feedback effects in non-mammalian species (though with different details) led Roenneberg, Daan and Merrow (2003, p. 190) to conclude that “the clock changes the properties of the input.”

We turn briefly to the output pathway. Recent research employing micro-arrays has resulted in a recognition that there are hundreds of genes whose expression is clock controlled—in some cases 10% of those studied (Panda et al., 2002). Some of the cycling proteins figure in rate regulating steps in reaction pathways, indicating circadian regulation of key metabolic functions. One example receiving attention recently involves clock regulation of the mitochondrial oxidative phosphorylation pathway that alters the rate of oxidative-reduction (redox) reactions within cells. Among the affected molecules are crucial oscillator components Clock, BMAL1, and NPAS2 (Rutter, Reick, Wu, & McKnight, 2001). As Panda and Hogenesch (2004, p. 381) concluded, “these observations may indicate that genes involved in redox regulation are both outputs of the clock as well as feedback on clock function.” This is just one illustration of how investigation of timekeeping began reductively, with a focus on molecular interactions, but later incorporated causal interactions among higher level systems.

Memory. In a similar manner, research on memory consolidation has ascended from obtaining an intracellular account back to system-level processing. By showing how the responsiveness of an individual neuron can be altered, LTP explains one aspect of learning. But the specific change in neural connectivity that occurs on a given learning trial depends upon what has already been learned. In particular, if the stimulus is very similar to others for which the desired response has already been learned, then it is appropriate to strengthen the connections already involved. If, on the other hand, the desired response is different than that learned for these other stimuli, then it is necessarily to learn a new response specific to that stimulus. In studying LTP, investigators typically focused on just one part of the hippocampus (e.g., CA1), but Rolls and Treves (1998) proposed that the distinctive organization of the hippocampus is crucial to how learning will occur. In particular, the various groups of cells comprising the hippocampus are organized into a loop.

Starting from the entorhinal cortex (EC), two pathways lead to the CA3, one direct and one mediated by the dentate gyrus (DG). The connectivity pattern in the DG ensures that only a few cells fire in response to a given input, thereby separating signals. In contrast, the CA3 cells have many recurrent projections back onto themselves, enhancing their ability to generate a similar overall response to a range of similar inputs. Both kinds of processing contribute to the input to CA1, the area in which LTP has been most studied. Moreover, the CA1 cells send a signal back to the EC, which has broad connections to other areas of cortex. Since the anterograde amnesia resulting from destruction of the hippocampus is time-delimited, leaving older memories intact, it is assumed that long-term episodic memories are ultimately stored in cortex. The output pathways from CA1 through EC to cerebral cortex have been proposed to figure in the ability of the hippocampus to train cells in the cerebral cortex that ultimately store the memories (McClelland, McNaughton, & O'Reilly, 1995). Thus, memory consolidation involves causal relations between brain regions, not just LTP within a region.

Other research also points to the importance of relating the cell and molecular research on LTP to the system level. It has often been assumed that memory consolidation involves a one-time process of changing connections between cells. Yet, animal learning studies in the 1960s produced extensive evidence that after learned behavior was elicited, it had to be reconsolidated or it was subject to disruption by electroshock or other treatment that impairs learning (Lewis, Bregman, & Mahan, 1972). These results were largely ignored by researchers investigating LTP until Sara (2000) encountered the phenomenon while investigating the effects of a NMDA receptor antagonist on rats performing a well-learned maze task. The rats later showed amnesia for the task, indicating that the same consolidation processes required in initial learning were required to keep the memory after recall. Much of the subsequent research has focused on protein synthesis inhibitors, as suggested by the role of protein synthesis in LTP. But what is important for our purposes is that frequently different parts of the brain than the hippocampus are crucial for reconsolidation of memories (Taubenfeld, Milekic, Monti, & Alberini, 2001). This is prompting researchers to explore interactions between numerous brain areas in the process of memory consolidation and reconsolidation.

Conclusions

Reductionistic research is often viewed as resulting in lower-level accounts that supplant the need for inquiry at higher levels. We have offered two cases to bolster our contention that this is not an accurate portrayal of even highly successful reductionistic research. Research on memory consolidation and circadian rhythms has been impressive in the detail it has provided on the genetic and biochemical processes involved in parts of the systems responsible for the phenomena of interest. But, far from supplanting the need for higher-level accounts, they have themselves

pointed to the need for higher-level accounts that capture how the systems in question are organized and how they interact with other systems. These interactions are at a level of organization above that targeted in reductionist research.

When seen in the context of the new mechanistic philosophy of science, the need for lower-level accounts to be supplemented with information from higher levels is to be expected. Lower-level inquiries investigate the parts and operations of a mechanism, and often the parts and operations within these parts. But parts and operations in a mechanism must be organized and orchestrated to produce the phenomenon of interest, and knowledge about the parts and operations alone does not provide information about organization. Moreover, the mechanism is situated in a context, and its behavior is influenced by the other occupants of that context. This too cannot be inferred from knowledge of the parts and their operations.

The entrée that cognitive neuroscience provides for linking cognitive operations to the brain has not been welcomed by all cognitive scientists, especially insofar as that opens the prospect of reductionist inquiry that attempts to dispense with cognitive explanations in terms of molecular ones. But that prospect, if our two case examples are an indication, is unlikely (which is not to say that some researchers may focus their entire attention on molecular mechanisms). Rather, the more that is learned about the parts and operations figuring in a mechanism, the more it becomes important to understand how the mechanism is organized and how it is situated in a broader context where it both affects, and is affected by, other entities. Cognitive scientists can draw upon the fruits of reductionist inquiry without losing their distinctive focus on understanding cognitive activities at a relatively high level of organization.

References

- Agranoff, B. W., Davis, R. E., & Brink, J. J. (1966). Chemical studies on memory fixation in goldfish. *Brain Research*, 1, 303-309.
- Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. London: Routledge.
- Bechtel, W., & Richardson, R. C. (1993). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Princeton, NJ: Princeton University Press.
- Bickle, J. (2003). *Philosophy and neuroscience: A ruthlessly reductive account*. Dordrecht: Kluwer.
- Bliss, T. V. P., & Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology*, 232, 331-356.
- Burnham, W. H. (1903). Retroactive amnesia: Illustrative cases and a tentative explanation. *American Journal of Psychology*, 14, 382-396.
- Cartwright, N. (1999). *The dappled world: A study of the boundaries of science*. Cambridge: Cambridge University Press.
- Craver, C. (2003). The making of a memory mechanism. *Journal of the History of Biology*, 36, 153-195.
- 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.
- Kim, J. (1998). *Mind in a physical world*. Cambridge, MA: MIT Press.
- Konopka, R. J., & Benzer, S. (1971). Clock mutants of *Drosophila melanogaster*. *Proceedings of the National Academy of Sciences (USA)*, 89, 2112-2116.
- Lewis, D. J., Bregman, N. J., & Mahan, J. (1972). Cue-dependent amnesia in rats. *Journal of Comparative and Physiological Psychology*, 81, 243-247.
- Machamer, P., Darden, L., & Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science*, 67, 1-25.
- McClelland, J. L., McNaughton, B., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3), 419-457.
- Panda, S., Antoch, M. P., Miller, B. H., Su, A. I., Schook, A. B., Straume, M., et al. (2002). Coordinated transcription of key pathways in the mouse by the circadian clock. *Cell*, 109(3), 307-320.
- Panda, S., & Hogenesch, J. B. (2004). It's all in the timing: many clocks, many outputs. *Journal of Biological Rhythms*, 19(5), 374-387.
- Pittendrigh, C. S. (1960). Circadian rhythms and the circadian organization of living systems. *Cold Spring Harbor Symposia on Quantitative Biology*, 25, 159-184.
- Roenneberg, T., Daan, S., & Mrosovsky, M. (2003). The art of entrainment. *Journal of Biological Rhythms*, 18(3), 183-194.
- Rollag, M. D., Berson, D. M., & Provencio, I. (2003). Melanopsin, Ganglion-Cell Photoreceptors, and Mammalian Photoentrainment. *J Biol Rhythms*, 18(3), 227-234.
- Rolls, E. T., & Treves, A. (1998). *Neural networks and brain function*. Oxford: Oxford University Press.
- Rutter, J., Reick, M., Wu, L. C., & McKnight, S. L. (2001). Regulation of Clock and NPAS2 DNA binding by the redox state of NAD cofactors. *Science*, 293(5529), 510-514.
- Sara, S. J. (2000). Retrieval and reconsolidation: Toward a neurobiology of remembering. *Learning and Memory*, 7, 73-84.
- Schouten, M., & Looren de Jong, H. (Eds.). (2007). *The matter of the mind: Philosophical essays on psychology, neuroscience and reduction*. Oxford: Blackwell.
- Taubenfeld, S. M., Milekic, M. H., Monti, B., & Alberini, C. M. (2001). The consolidation of new but not reactivated memory requires hippocampal C/EBP β . *Nature Reviews Neuroscience*, 4, 813-818.
- Thagard, P. (2006). *Hot thought: Mechanisms and applications of emotional cognition*. Cambridge, MA: MIT Press.
- Welsh, D. K., Logothetis, D. E., Meister, M., & Reppert, S. M. (1995). Individual neurons dissociated from rat suprachiasmatic nucleus express independently phased circadian firing rhythms. *Neuron*, 14(4), 697-706.