

Understanding the Brain as an Endogenously Active Mechanism

William Bechtel (bill@mechanism.ucsd.edu)

Department of Philosophy, University of California, San Diego
La Jolla, CA 92093-0119 USA

Adele Abrahamsen (aabrahamsen@ucsd.edu)

Center for Research in Language, University of California, San Diego
La Jolla, CA 92093 USA

Abstract

Although a reactive framework has long been dominant in cognitive science and neuroscience, an alternative framework emphasizing dynamics and endogenous activity has recently gained prominence. We review some of the evidence for endogenous activity and consider the implications not only for understanding cognition but also for accounts of explanation offered by philosophers of science. Our recent characterization of *dynamic mechanistic explanation* emphasizes the coordination of accounts of mechanisms that identify parts and operations with computational models of their activity. These can, and should, be extended to incorporate attention to mechanisms that are not only active, but endogenously active.

Keywords: philosophy of science; mechanistic explanation; dynamics; endogenous brain activity, resting state fMRI, brain default network

Introduction

Observe a living organism, from a bacterium to a fellow human being, and you see an endogenously active system. Introspect and you will observe, as did William James, a continual flow of thoughts. If pressed, most cognitive scientists will acknowledge that neural systems—from individual neurons to the brain as a whole—exhibit endogenous activity. That is, some of the activity is internally (Greek *endo*) produced (German *gennan*); the causes and control of this activity is inside the system rather than reactive to inputs from outside the system. But cognitive scientists tend to disregard this when designing studies. Those in psychology present discrete stimuli in structured tasks designed to permit statistical analysis of the behavioral effects of independent variables. Those in neuroscience, following the tradition of Charles Scott Sherrington (1923), commonly treat the brain as a reactive system in which sensory inputs initiate neural processing that results ultimately in motor responses. They may stimulate specific neurons or provide sensory inputs with specific properties so that recorded neural activity can be analyzed in terms of responses to inputs. In both fields, variations in activity that cannot be associated with an input are treated as random fluctuations (noise). There is no doubt that this reactive framework in psychology and neuroscience has been enormously productive in identifying the parts, operations, and organization of the mechanisms responsible for cognition. It soon reaches its limits, though, in seeking accounts of the orchestrated functioning of those components: their dynamics and coordination in real time.

The investigation of endogenous activity, though less influential, has historical roots nearly as deep as those of the reactive approach. It was promoted by Thomas Graham Brown (1914), for example, who studied decerebrate and deafferented cats in Sherrington's laboratory at Liverpool from 1910 to 1913. He found that the isolated spinal cord, even when not receiving inputs, generates patterns of activity comparable to those exhibited during motor behavior elicited by stimuli. Brown's emphasis on endogenous activity initially was largely ignored (for discussion, see Stuart & Hultborn, 2008) but was revived several decades later when biologists recognized a class of neural circuits—*central pattern generators*—whose self-sustaining patterns of activity generated rhythmic motor behavior even in the absence of sensory input. After Wilson and Wyman (1965) pioneered this construct in their account of locust flight, others identified central pattern generators in the brain stem and spinal cord for walking, swimming, respiration, circulation, and other behaviors for which oscillatory control was crucial (Grillner, 2003). Endogenous activity has received far less attention from those studying sensory processing and central cognition rather than motor control, despite indications of endogenous oscillatory activity in cerebral cortex using techniques ranging from single cell recording to EEG and fMRI. In the next section we describe highlights from this research and in the subsequent section briefly explore the implications for reconstruing how we understand cognitive activity. Most important, if the conception of the brain as endogenously active is taken seriously, it profoundly challenges the reactive perspective that has dominated much of cognitive science as well as neuroscience: stimuli or tasks must be regarded not as initiating activity in an inactive system, but rather as perturbing endogenous dynamic behavior.

The slow pace at which these fields are achieving a change of perspective is unsurprising considering the history of other sciences. Although Max Planck was exaggerating when he said "A new scientific truth does not triumph by convincing its opponents . . . but rather because its opponents eventually die . . .," the considerable costs and uncertain benefits of change make it a tough sell. Uneven acceptance of Einstein's revolutionary proposals is a familiar example. Less remarked upon is the delayed impact of changes in the sciences on *philosophy* of science. For example, this young field (which did not even have a journal until 1934) did not exhibit acute concern with the epistemological foundations of science until it was confronted with Ein-

stein's proposals and their aftermath—a response that necessarily involved at least a short delay. However, delays in uptake have been far greater for developments in sciences other than physics, notably the biological and cognitive sciences. Philosophers of science did not even recognize the dominant mode of explanation in these sciences—mechanistic explanation—until the 1990s and especially after 2000. More recently, we have argued that such developments as computational modeling of the dynamics of cognitive and neural mechanisms require philosophers of science to extend their notion of mechanism to include *dynamic mechanistic explanation*. In the last section of this paper we will briefly characterize these two explanatory frameworks and consider how the philosophical understanding of dynamic mechanistic explanation can incorporate the implications of scientific work on endogenous activity.

Evidence that the Brain is Endogenously Active

Although lesion and stimulation techniques have been important in identifying brain regions involved in different cognitive activities, since the mid-20th century the greatest insights have come from techniques in which researchers record brain activity of individual neurons (single or multi-cell recording) or brain regions (EEG and fMRI). Most commonly these techniques have been employed within the reactive framework in which stimuli are presented or tasks are assigned, responses within the brain recorded, and these responses pooled for analysis to remove variability not associated with the intervention.

Each of these techniques, though, also has been employed in ways that reveal endogenous brain activity. Notably, Rodolfo Llinás employed intracellular recordings to identify systematic variations in the conductance of calcium ions across neural membranes. He showed how the manner in which these conductances varied through time enabled neurons in the inferior olive, a brainstem nucleus, to function as single-cell oscillators “capable of self-sustained rhythmic firing independent of synaptic input” (Llinás, 1988, p. 1659). (For a review of evidence and models showing how these intrinsic oscillations when combined with synaptic processes can generate synchronous thalamocortical oscillations, see Destexhe & Sejnowski, 2003.)

A second line of evidence for endogenous brain activity, consistent with that of single-cell recording, emerged from earlier studies by Hans Berger (1929) pioneering the identification of distinctive waveforms in electroencephalograph (EEG) recordings of brain activity. When he presented no stimuli or task demands but simply had subjects sit awake with their eyes closed, he obtained high-amplitude oscillations between 8 and 12 Hz that he dubbed *alpha waves*. When subjects instead viewed a stimulus or solved a problem, alpha waves were supplanted by lower-amplitude, higher-frequency *beta waves* (12-30 Hz). Soon thereafter it was determined that the EEG signal captured, not action potentials, but rather synchronized sub-threshold electrical potentials across a population of neurons. In the 1960s, the

development of digital EEG and of powerful statistical techniques for decomposing complex EEG signals into component waveforms brought further discoveries; notably, very high-frequency (25-100 Hz) *gamma waves* were prominent in addition to beta waves when people performed various cognitive tasks. Moreover, synchronized oscillations at all of these frequencies were found in both active and passive conditions, but at different amplitudes.

Thus, both single-cell recording and EEG studies have provided evidence for endogenous brain activity. In this paper we will focus on yet another line of evidence offered by recent work on *resting-state fMRI*. The BOLD (blood oxygen level dependent) signal employed in fMRI research registers the oxygen concentrations in the brain within areas that can be as small as 2 mm. Until recently fMRI research focused nearly exclusively on finding higher values in the BOLD signal when a task condition is compared to a control or resting state condition.¹ For example, semantic processing of words (task condition) would be contrasted to reading words aloud (control condition) or to lying still in the scanner with eyes closed (resting condition). The interest in neuroimaging during a resting state, rather than during task performance, developed from researchers' occasional observations that a number of brain areas routinely exhibited less activity in task situations than in the resting state. To explore further these intriguing observations, Shulman et al. (1997) conducted a meta-analysis of studies in which a task condition was compared to a non-task condition in which the same stimulus was present. They found that the areas commonly less active in task situations included posterior cingulate cortex (PCC), precuneus, inferior parietal cortex (IPC), left dorsal lateral prefrontal cortex (left DLPFC), and a medial frontal strip that continued through the inferior anterior cingulate cortex (ACC), left inferior frontal cortex, and left inferior frontal gyrus to the right amygdala. Turning the focus from the fact that these areas are less active during tasks to the fact that they are more active in the absence of task requirements, Raichle and his collaborators (Raichle et al., 2001) suggested that together these areas constitute a *default network*.

A major advance in understanding the default network resulted from analyzing the temporal dynamics of the BOLD signal. A pioneering dynamical analysis of fMRI data was provided by Biswal, Yetkin, Haughton, and Hyde (1995), who obtained BOLD signal values every 250 msec after a hand movement and identified spontaneous low frequency

¹ In referring to resting states, the assumption is not that the subject's brain is resting, but that he or she is not engaged in a specific task or responding to a specific stimulus. Often the subject is asked to fixate on a cross-hair or lie still in the scanner with eyes closed but not asleep. Fluctuations in activity that can be linked to physiological activity (cardiac or respiratory activity) are eliminated from the data through linear regression. In a critique of this research, Morcom and Fletcher (2007) focused on the privileging of the resting state. The insights into the default network on which we focus, however, do not rely on the resting state being privileged but simply as revealing ongoing activity in brain networks not employed in cognitive tasks.

(less than 0.1 Hz) fluctuations in sensorimotor cortex. These fluctuations were synchronized across the left and right hemispheres and with those in other motor areas, which was interpreted as evidence of functional connectivity among all these areas. Accordingly, the approach is referred to as *functional connectivity MRI* (fcMRI).

Employing fcMRI, Greicius, Krasnow, Reiss, and Menon (2003) demonstrated that if they used the PCC as a seed for statistical analysis, they could identify synchronized fluctuations in a large cluster of areas: medial prefrontal cortex (including inferior ACC and orbitofrontal cortex), left DLPFC, inferior parietal cortex bilaterally, left inferolateral temporal cortex, and left parahippocampal gyrus. Taking instead the inferior ACC as the seed area, they found correlated fluctuations in the PCC, medial prefrontal cortex/orbital frontal cortex, the nucleus accumbens, and the hypothalamus/midbrain. Since these regions were virtually the same as those showing activity in Shulman's resting state data, Greicius et al. construed this as evidence for "a cohesive, tonically active, default mode network" (p. 256) with two subnetworks.

While the default network exhibits greater activity in the resting state than in task conditions, the areas showing greater activity in task conditions still generate a BOLD signal in the resting state and one can find correlations in the dynamics across these areas (synchronized oscillations). These synchronized oscillations are, however, out of phase with those in the default network. Comparing the default network with one that exhibited greater activation in an attention-demanding task (intraparietal sulcus, frontal eye field, middle temporal region, supplementary motor areas, and the insula), Fox et al. (2005) described oscillations in the two networks as anticorrelated, whereas oscillations for different areas within each network were positively correlated. This shows that both the default network and the network involved in attention-demanding tasks are coordinating their activities within themselves in the absence of external stimulation or task demands.

Researchers subsequently identified additional networks using this strategy. That is, a set of areas with correlated dynamics (synchronized oscillations) under resting state conditions were posited to constitute a network, further evidenced by negative correlations with other networks (e.g., Mantini, Perrucci, Del Gratta, Romani, & Corbetta, 2007, differentiate six anticorrelated networks). Fox and Raichle (2007) concluded: "A consistent finding is that regions with similar functionality—that is, regions that are similarly modulated by various task paradigms—tend to be correlated in their spontaneous BOLD activity."

Although the oscillations revealed in fMRI are of a much lower frequency (< 0.1 Hz) than those usually reported in EEG (1-80 Hz), researchers have found ways to relate them. Mantini et al., for example, found that "Each brain network was associated with a specific combination of EEG rhythms, a neurophysiological signature that constitutes a baseline for evaluating changes in oscillatory signals during active behavior" (p. 13170). For example, the default net-

work showed positive correlations with amplitude in alpha and beta band oscillations while the attention network exhibited negative correlations in these frequency bands. These correlations may reflect systemic coherence in brain functioning. In the cortex of mammals, the amplitude (power density) of EEG oscillations has been found to be inversely proportional to their frequency ($1/f$). Even more interesting, the phase of lower-frequency oscillations seems to modulate the amplitude of those at higher frequencies, which results in a nesting relation between the frequency bands. (Lakatos et al., 2005, refer to this as "oscillatory hierarchy hypothesis") In addition, oscillations at lower frequencies tend to synchronize over more widely distributed areas of the brain than those at higher frequencies (Buzsáki & Draguhn, 2004). Such coupling can be particularly important when the brain is perturbed by a stimulus, since a modulation in low-frequency oscillations can, through phase-locking with higher-frequency oscillations, yield rapid changes at those frequencies.

The Significance of Endogenous Brain Activity for Understanding Cognition

One might acknowledge endogenous activity in various brain networks, but deny that it is of any cognitive significance. Perhaps it merely reflects basic metabolic activity and bears no implications for cognition. However, the fact that each network oscillates at a characteristic frequency, rather than fluctuating randomly, suggests that endogenous activity has implications for understanding brain activity generally—including activity during cognitive functioning. We briefly explore different ways in which endogenous activity may be important for understanding the brain as a system for cognition.

First, if a mechanism responds to a stimulus by increasing its activity, and that activity already is oscillating, response to the stimulus will vary depending on the phase of the oscillation when the stimulus arrives. This is true of individual neurons. If the membrane voltage of a neuron oscillates endogenously in a range below zero mV, as the evidence developed by Llinás and others indicates, then it will require stronger input to exceed the threshold for generating an action potential when it happens to be at its most negative phase. The same principle applies to populations of neurons whose oscillations are synchronized. In a variety of tasks in which a stimulus evokes a behavioral response, it is known that the response correlates with the magnitude of the BOLD signal. Fox, Snyder, Zacks, and Raichle (2005) therefore investigated whether these effects could be explained by synchronized spontaneous fluctuations in neuronal activity detectable with fMRI. Subjects were instructed to press a button with the right hand when a stimulus was detected, resulting in evoked activity in the left somatosensory cortex. The researchers hypothesized that the ongoing spontaneous fluctuations in the right somatosensory cortex provided an accurate measure of the spontaneous contribution to activity in the left somatosensory area at each timestep and succeeded in showing that these sponta-

neous fluctuations contributed significantly to the amplitude of blood flow in the left somatosensory areas after each stimulus. In fact, the task-related increased blood flow could be analyzed as a linear addition to the current amplitude of the spontaneous fluctuation. From this they inferred that the underlying spontaneous fluctuations affected perception and behavior. They supported this conclusion more directly in a subsequent study, in which they determined that spontaneous fluctuations accounted for variability in the force with which subjects pressed the button (Fox, Snyder, Vincent, & Raichle, 2007). When subjects were instructed as to how forcefully they should press the button, the pattern of neuronal activity was very different than that which arose when they were not instructed, allowing the investigators to discount the possibility that what they took to be spontaneous variability was in fact an evoked response. Thus, their study can be taken as initial evidence that the variability in endogenous brain activity is one source of the variability in measures of cognitive activity.

Second, endogenous activity in the brain's default network is the most obvious candidate for the neural underpinnings of *mindwandering* (Antrobus, Singer, Goldstein, & Fortgang, 1970). In one of the early fMRI studies using the resting state, Andreasen et al. (1995) queried subjects about what they were doing and elicited reports of being engaged in "a mixture of freely wandering past recollection, future plans, and other personal thoughts and experiences." Since these activities involve episodic memory, and episodic memory tasks are among those which do not lead to lower activity in the default network, Andreasen et al. and subsequent researchers (e.g., Buckner & Carroll, 2007) have suggested that the default network is involved in recalling personal experiences and anticipating future ones. Intriguingly, Li, Yan, Bergquist, and Sinha (2007) correlated trials on which subjects failed to detect stop signals in behavioral tasks with increased activity in the default network, as one would expect if that network were involved in a person thinking distracting thoughts about past and future experiences. One factor that renders problematic such a characterization of the activity of the default network is that the oscillatory behavior of the default network is maintained as well in sleep (Fukunaga et al., 2006) and under anesthesia (Vincent et al., 2007), when presumably spontaneous thoughts are not occurring.

Third, endogenous brain activity might be crucial for building and maintaining certain types of organization in the nervous system required for cognitive activity. There is growing evidence that the brain exhibits *small-world* organization (Watts & Strogatz, 1998) in which most connections link neighboring neurons, creating clusters that can collaborate in processing specific information, but a few long range connections enable overall coordination (Sporns & Zwi, 2004). There also is evidence that while most brain areas have connections to only a few other areas, some have a large number of connections, thereby constituting hubs. Such an architecture provides a highly efficient organization for information processing, and it is notable that the default

network itself exhibits a small-world architecture with hubs. An important question is how such organization might arise. Rubinov, Sporns, van Leeuwen, and Breakspear (2009) advanced the intriguing possibility that oscillatory neurons, developing connections when synchronized, might self organize into a small world network with hubs. In support of this proposal they described a model by Gong and van Leeuwen (2004) that employs a logistic map activation function for individual units that endogenously exhibit chaotic behavior. This enables the emergence of temporary patterns of synchronized oscillations even in the absence of external stimulation. A Hebbian learning procedure establishes new connections between pairs of units whose activity is synchronized and prunes those between unsynchronized units. Even when these networks begin with random connectivity, they develop clusters linked to each other through hubs. However, in real brains the initial state already involves local regions with interconnections and experience further shapes the emerging organization such that the outcome is a highly correlated brain capable of maintaining multiple anticorrelated networks. That is, the architecture of the information processing system may be shaped by both endogenous and exogenous activity.

In this section we have considered three suggestions as to how endogenous activity in the brain may contribute to its functioning as a cognitive system. Although it is too early to judge which will prove most fruitful, clearly the time for dismissing endogenous activity as mere noise has passed.

Endogenously Activity and Mechanistic Explanation

The evidence for endogenous activity in brains presents challenges not only to the ways in which cognitive scientists understand cognitive activity but also to philosophers' construal of the explanatory frameworks used in science. We mentioned above that these construals lag behind the sciences, often far more than necessary. Until recently, philosophical accounts of explanation focused primarily on laws and construed explanation as the subsumption of phenomena to be explained under these laws. While such an approach might work in physics, where there are many well established laws, it does not characterize explanations in the life sciences, where there are few laws but an abundance of phenomena to be explained (Cummins, 2000). What form of explanation is appropriate? In the past 20 years a number of philosophers of science have finally paid attention to biologists and, following their lead, construed explanation as the characterization of the mechanism responsible for a phenomenon of interest (Bechtel & Richardson, 1993; Bechtel & Abrahamsen, 2005; Machamer, Darden, & Craver, 2000; Thagard, 2006).

Although there are minor differences among these various accounts of mechanistic explanation, they concur in construing a mechanism as consisting of component parts, each of which performs one or more operations. Each operation produces change in another part that triggers or affects the operation of that part, and so forth. Cognitive psychologists,

traditionally have posited operations that transform, copy, or move representations without localizing them in parts of the brain. Cognitive neuroscientists (and growing numbers of cognitive psychologists) emphasize localization and choose operations at the appropriate grain for their brain recording technology (Bechtel, 2008).

Given the focus on specifying a mechanism to explain a given phenomenon, it is natural to conceive of the mechanism as having a specific beginning condition and continuing its operations until its task is completed. This sequential conception of mechanism is most clearly captured in the definition offered by Machamer, Darden, and Craver (2000): "Mechanisms are entities [parts] and activities [operations] organized such that they are productive of regular changes from start or set-up to finish or termination conditions." If the start or set up conditions involve a stimulus or task originating from outside the mechanism, we arrive at the construal of a mechanism not only as sequential but also as reactive.

This reactive conception of a mechanism accords well with the accounts offered in many areas of biology and cognitive science, but it is not adequate to characterize endogenously active systems as discussed in the previous sections. A sequentially organized mechanism will not exhibit endogenous activity. A minimal first step towards a mechanism capable of endogenous activity retains the general sequential conception of the overall functioning of the mechanism but allows operations that are viewed as later in the sequential order to feed back, either negatively or positively, on operations thought of as earlier. With even a single negative feedback loop it is possible to generate oscillatory behavior. It has long been known (Goodwin, 1965) that if the operations are appropriately non-linear *and* the system is open to sources of energy, these oscillations may be self-sustained and not dampen to a steady state over time. The same is true of mechanisms employing positive feedback or cyclic organization (see Bechtel & Abrahamsen, 2010).

Accommodating these organizational principles requires dropping the sequential characterization of a mechanism and instead coordinating accounts of parts and operations with accounts of their dynamics. The conception of mechanism hence becomes more dynamic: "A mechanism is a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism, **manifested in patterns of change over time in properties of its parts and operations**, is responsible for one or more phenomena" (Bechtel & Abrahamsen, in press). Accounts that utilize this conception exemplify what we have recently called *dynamic mechanistic explanation*. Often such accounts incorporate computational modeling of the real-time dynamics produced by feedback loops and other forms of cyclic organization. Moreover, a dynamic conception of mechanism and mechanistic explanation is compatible with the non-sequential organization, non-linear interactions, and openness to energy required for endogenous operation.

A self-sustaining oscillatory mechanism can account for the endogenous activity found in the brain, but now new explanatory tasks arise. First, the phenomenon of interest is typically not generated by a single oscillatory mechanism but by the coordinated behavior of multiple oscillators. Since Huygens we have known that if a signal can be passed between oscillators, they can synchronize their oscillations. However, depending on the particular ways in which oscillators are organized into a system, a population of oscillators can come to exhibit extremely complex behavior. Second, even a single oscillator can be perturbed by external inputs and the resulting change in its functioning can be complex. Complexity is even greater when a population of oscillators already exhibiting complex behavior is perturbed. These are the sorts of challenges faced in understanding how the brain, viewed as an endogenously active system, is presented with stimuli or tasks. Philosophical accounts of explanation must also reflect these challenges confronted in neuroscience and cognitive science.

References

- Andreasen, N. C., O'Leary, D. S., Cizadlo, T., Arndt, S., Rezai, K., Watkins, G. L., et al. (1995). Remembering the past: two facets of episodic memory explored with positron emission tomography. *American Journal of Psychiatry*, 152 (11), 1576-1585.
- Antrobus, J. S., Singer, J. L., Goldstein, S., & Fortgang, M. (1970). Mindwandering and cognitive structure. *Transactions of the New York Academy of Sciences*, 32, 242-252.
- Bechtel, W. (2008). *Mental mechanisms*. London: Routledge.
- 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). Complex biological mechanisms: Cyclic, oscillatory, and autonomous. In C. A. Hooker (Ed.), *Philosophy of complex systems. Handbook of the philosophy of science, Volume 10*. New York: Elsevier.
- Bechtel, W., & Abrahamsen, A. (in press). Dynamic mechanistic explanation: Computational modeling of circadian rhythms as an exemplar for cognitive science. *Studies in History and Philosophy of Science Part A*.
- Bechtel, W., & Richardson, R. C. (1993). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Princeton, NJ: Princeton University Press.
- Berger, H. (1929). Über daas Elektroenkephalogramm des Menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, 87 (527-570).
- Biswal, B., Yetkin, F. Z., Haughton, V. M., & Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine*, 34 (4), 537-541.
- Brown, T. G. (1914). On the nature of the fundamental activity of the nervous centres; together with an analysis of the conditioning of rhythmic activity in progression, and a

- theory of the evolution of function in the nervous system. *The Journal of Physiology*, 48 (1), 18-46.
- Buckner, R. L., & Carroll, D. C. (2007). Self-projection and the brain. *Trends in Cognitive Sciences*, 11 (2), 49-57.
- Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, 304 (5679), 1926-1929.
- Cummins, R. (2000). "How does it work?" versus "what are the laws?": Two conceptions of psychological explanation. In F. Keil & R. Wilson (Eds.), *Explanation and cognition* (pp. 117-144). Cambridge, MA: MIT Press.
- Destexhe, A., & Sejnowski, T. J. (2003). Interactions between membrane conductances underlying thalamocortical slow-wave oscillations. *Physiological Reviews*, 83 (4), 1401-1453.
- Fox, M. D., & Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, 8 (9), 700-711.
- Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the United States of America*, 102 (27), 9673-9678.
- Fox, M. D., Snyder, A. Z., Vincent, J. L., & Raichle, M. E. (2007). Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron*, 56 (1), 171-184.
- Fox, M. D., Snyder, A. Z., Zacks, J. M., & Raichle, M. E. (2005). Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses. *Nature Neuroscience*, 9, 23-25.
- Fukunaga, M., Horovitz, S. G., van Gelderen, P., de Zwart, J. A., Jansma, J. M., Ikonomidou, V. N., et al. (2006). Large-amplitude, spatially correlated fluctuations in BOLD fMRI signals during extended rest and early sleep stages. *Magnetic Resonance Imaging*, 24 (8), 979-992.
- Gong, P., & van Leeuwen, C. (2004). Evolution to a small-world network with chaotic units. *Europhysics Letters*, 67, 328-333.
- Goodwin, B. C. (1965). Oscillatory behavior in enzymatic control processes. *Advances in Enzyme Regulation*, 3, 425-428.
- Greicius, M. D., Krasnow, B., Reiss, A. L., & Menon, V. (2003). Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences of the United States of America*, 100 (1), 253-258.
- Grillner, S. (2003). The motor infrastructure: from ion channels to neuronal networks. *Nature Reviews Neuroscience*, 4 (7), 573-586.
- Lakatos, P., Shah, A. S., Knuth, K. H., Ulbert, I., Karmos, G., & Schroeder, C. E. (2005). An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex. *Journal of Neurophysiology*, 94 (3), 1904-1911.
- Li, C.-S. R., Yan, P., Bergquist, K. L., & Sinha, R. (2007). Greater activation of the "default" brain regions predicts stop signal errors. *Neuroimage*, 38 (3), 640-648.
- Llinás, R. R. (1988). The intrinsic electrophysiological properties of mammalian neurons: Insights into central nervous system function. *Science*, 242 (4886), 1654-1664.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67, 1-25.
- Mantini, D., Perrucci, M. G., Del Gratta, C., Romani, G. L., & Corbetta, M. (2007). Electrophysiological signatures of resting state networks in the human brain. *Proceedings of the National Academy of Sciences*, 104 (32), 13170-13175.
- Morcom, A. M., & Fletcher, P. C. (2007). Does the brain have a baseline? Why we should be resisting a rest. *Neuroimage*, 37 (4), 1073-1082.
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences of the United States of America*, 98 (2), 676-682.
- Rubinov, M., Sporns, O., van Leeuwen, C., & Breakspear, M. (2009). Symbiotic relationship between brain structure and dynamics. *BMC Neuroscience*, 10 (1), 55.
- Sherrington, C. S. (1923). *The integrative action of the nervous system*. New Haven: Yale University Press.
- Shulman, G. L., Fiez, J. A., Corbetta, M., Buckner, R. L., Miezin, F. M., Raichle, M. E., et al. (1997). Common blood flow changes across visual tasks. II. Decreases in cerebral cortex. *Journal of Cognitive Neuroscience*, 9 (5), 648-663.
- Sporns, O., & Zwi, J. D. (2004). The small world of the cerebral cortex. *Neuroinformatics*, 2 (2), 145-162.
- Stuart, D. G., & Hultborn, H. (2008). Thomas Graham Brown (1882-1965), Anders Lundberg (1920-), and the neural control of stepping. *Brain Research Reviews*, 59 (1), 74-95.
- Thagard, P. (2006). *Hot thought: Mechanisms and applications of emotional cognition*. Cambridge, MA: MIT Press.
- Vincent, J. L., Patel, G. H., Fox, M. D., Snyder, A. Z., Baker, J. T., Van Essen, D. C., et al. (2007). Intrinsic functional architecture in the anaesthetized monkey brain. *Nature*, 447 (7140), 83-86.
- Watts, D., & Strogatz, S. (1998). Collective dynamics of small worlds. *Nature*, 393 (440-442).
- Wilson, D. M., & Wyman, R. J. (1965). Motor output patterns during random and rhythmic stimulation of locust thoracic ganglia. *Biophysical Journal*, 5 (2), 121-143.