

After the Philosophy of Mind: Replacing Scholasticism with Science*

Anthony Chemero and Michael Silberstein^{†‡}

We provide a taxonomy of the two most important debates in the philosophy of the cognitive and neural sciences. The first debate is over methodological individualism: is the object of the cognitive and neural sciences the brain, the whole animal, or the animal—environment system? The second is over explanatory style: should explanation in cognitive and neural science be reductionist-mechanistic, interlevel mechanistic, or dynamical? After setting out the debates, we discuss the ways in which they are interconnected. Finally, we make some recommendations that we hope will help philosophers interested in the cognitive and neural sciences to avoid dead ends.

1. Introduction. The philosophy of mind is over. The two main debates in the philosophy of mind over the last few decades about the *essence* of mental states (are they physical, functional, phenomenal, etc.) and over mental content have run their course. Positions have hardened; objections are repeated; theoretical filigrees are attached. These relatively armchair discussions are being replaced by empirically oriented debates in philosophy of the cognitive and neural sciences. We applaud this, and agree with Quine that “philosophy of science is philosophy enough” (1966, 149). The purpose of this paper is first to provide a guide to philosophy of mind’s successor debates in philosophy of cognitive science, and second to suggest resolutions for some obvious potential conflicts so as to avoid the scholastic pitfalls that plagued philosophy of mind. We will discuss two not quite distinct debates: the first over the proper object of study for the psychological sciences; the second over the explanatory style of the cognitive and neural sciences. After describing the two debates and

*Received February 2007; revised July 2007.

†To contact the authors, please write to: Anthony Chemero, Franklin and Marshall College, Lancaster, PA 17604-3003; e-mail: tony.chemero@fandm.edu; Michael Silberstein, Elizabethtown College, One Alpha Drive, Elizabethtown, PA 17022-2298, or University of Maryland, College Park, MD 20742; e-mail: silbermd@etown.edu.

‡Thanks to Colin Klein and Bill Bechtel for discussion of these issues.

Philosophy of Science 75 (January 2008) pp. 1–27. 0031-8248/2008/7501-0001\$10.00
Copyright 2008 by the Philosophy of Science Association. All rights reserved.

setting out the possible positions on them, we will discuss some inter-connections between positions and then point to what we take to be some important issues for the philosophy of the cognitive and neural sciences in the next few years.

2. Debate #1: On Methodological Individualism. Ever since Putnam's declaration that meanings are not in the head, there has been a vigorous debate over content externalism: is it possible to specify the *content* of mental states based purely on facts about the brain of the thinker, or must one also take into account features of the thinker's current or past environment? More recently, however this debate has shifted to one over whether the *vehicle* or content bearer is confined to the brain of the thinker. The modern incarnation of this debate stemmed from work in connectionist networks by Rumelhart et al. (1986). The suggestion there was that the pattern-completing brain was only a proper part of the cognitive system, the rest of which was external to the thinker's body. Rumelhart et al.'s example was solving mathematical problems on a chalkboard. In such a case, it was argued that the cognitive system included the brain, the chalkboard and the act of writing on the board. Thus was the debate begun in earnest. Before describing the possible positions in the debate, it is worth pointing out that this is not a debate over metaphysics as such. The question is not whether the mind is identical with or confined to the brain. Instead, the issue of concern is a scientific one: What is the object that the mature cognitive sciences ought to study?

We represent the debate in the form of a decision tree. (See Figure 1.)

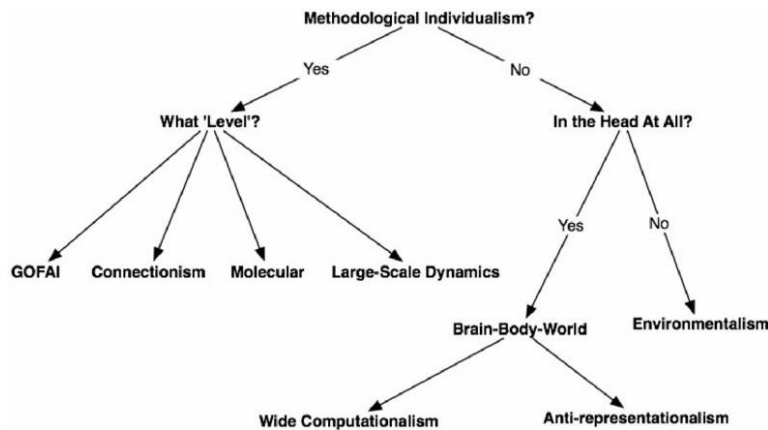


Figure 1. On methodological individualism.

The first cut on the debate is over whether cognitive systems are wholly confined to the heads of thinkers.

2.1. Cut 1: Is the Cognitive System All in the Head? Yes = Internalism. As noted above, one might hold that the object of psychological science is on the inside of the organism whether or not one believes that the contents of psychological states are determined by internal matters. This is a metaphysical issue that need not concern us here, however. If one believes that the cognitive system is confined to the head of the thinker, the next cut in this internalist branch of the philosophy of the cognitive sciences is over at what 'level' our explanations of cognitive systems ought to be pitched. We will return to this question shortly.

No = Externalism. Answering 'no' to this question is suggesting that cognitive systems are larger in spatial extent than brains. If one takes this externalist branch of the decision tree, the next pressing question is where in space cognitive systems are. We will return to this question after following the internalist side of the decision tree to the bottom.

2.2. Internalism Cut 2: At What 'Level' Should We Explain Cognitive Systems? For those who believe that cognitive systems are confined to the skull of the cognizer, the next question is at what 'level' the best explanations will be pitched. There are several possible answers to this question, the majority of which have been available for many years. First, one might believe that cognitive systems are computers, whose dynamics are to be explained in terms of symbolic representations and rules that determine their transformations. This sort of view is most strongly associated with Fodor (1975); Fodor and Pylyshyn (1988) call this the 'classical' view in cognitive science. Haugeland, who does not hold the view, calls it Good-Old-Fashioned-AI, or GOFAI (1985). A second possibility is that one might take the cognitive architecture to be connectionist networks. A view like this, championed by Churchland (1989) among many others, is pitched at a lower level than GOFAI. 'Connectionists' take it that the best explanations of cognitive systems will be subsymbolic and more closely connected to the actual activities of brain areas. The debate between GOFAI and connectionism has been raging for 20 years now, and some people apparently still find it worth thinking about. Of more obvious contemporary interest is a third possibility: one might go to a lower level still, and suggest that the best explanations of cognitive systems will be in terms of neurotransmitters and genetic activity. This view has been championed forcefully in recent years by Bickle (2003). Bickle argues for what he calls 'ruthless reductionism', according to which molecular neuroscience, along with appropriate bridge principles, will be able to account for all the laws and facts of psychology. A fourth possibility, one

that has gotten far too little attention from philosophers, is that the best way to explain cognition is in terms of large scale neural dynamics (see Cosmelli, Lachaux, and Thompson 2007 for a summary). Freeman has been putting forth explanations like this for many years (e.g., Skarda and Freeman 1987). More recently, Bressler and Kelso (2001) proposed similar explanations. Among philosophers, Thompson and Varela (2001) have endorsed this approach.

In addition to the level at which explanations are pitched, these views differ over the nature of mental representations. For GOFAI, representations are sentence-like; connectionists typically take mental representations to be distributed across clusters of simple processing units, and representable in terms of a vector in a n -dimensional state space, where n is the number of simple processing units. Some connectionists doubt the explanatory value of representational understandings of their networks (e.g., Ramsey 1997). In this respect, they are like Freeman and other proponents of a dynamical approach in neuroscience. Freeman has argued for years that it is a mistake to think that cognitive systems represent their environments (Freeman and Skarda 1990; Freeman 1999). It is also hard to imagine representations among neurotransmitters, though Bickle makes no commitment either way as to the value of representations in explaining cognition.

2.3. *Externalism Cut 2: Is the Cognitive System in the Head at All?*

No = Radical Environmentalism. The claim that the cognitive system is not in the head at all, that cognition is to be explained entirely in terms of the interactions of whole animals and their environments, may seem like an automatic nonstarter and an idea so crazy that no one would have held it. That is not so. Skinnerian behaviorists still make claims like this (Hineline 2006), and the later work of Gibson (1979) can be interpreted as making claims like this. In both cases, the claim is that all of the explanatory work can be done by carefully studying the ways active animals interact with the environment. In the Skinnerian case, one focuses on the subtle ways that animal behavior is shaped by environmental outcomes, and claims that reinforcement learning can account for the whole gamut of behavior. In the Gibsonian case, one focuses on the breathtaking amount of information available to a perceiver, especially one that is moving, and claims that this information is sufficient for perception of the environment without the addition of information stored in the brain. Note that neither Gibson nor Skinnerians claim that the brain is not importantly *involved in* cognition; rather they claim that psychologists can do all their explanatory work without referring to the brain. Both Skinnerians and Gibsonians have been very successful as psychologists, and both groups have achieved results that are undeniable psychological mile-

stones. So these views are neither nonstarters, nor obviously crazy. Nonetheless, radical environmentalism is at odds with most contemporary psychology.

Yes = Brain-Body-Environment (BBE) Interaction. To follow Rumelhart et al. (1986) is to take human competence at long division as encompassing brain, body, and portions of the environment. On this view, the brain's job is simple pattern completion, and long division is possible only because one can control one's hands with the chalk and then read one's inscriptions on the board. In cases like this, the board and the acts of writing on it and reading from it are part and parcel of the ability. Hence any explanation of the abilities of the cognitive system will have to include the brain, the writing, the board, and the reading, as well as the interactions among them. This sort of view has gotten a lot of attention in recent years under the names 'embodied cognition', 'situated cognition', and 'embedded cognition' (Brooks 1991; Varela, Thompson, and Rosch 1991; McClamrock 1995; Clark 1997, 2003; Anderson 2003; Wilson 2004; Wheeler 2005b). On this view, our bodies are well designed tools, making them easy for our brains to control. For example, our kneecaps limit the degrees of motion possible with our legs, making balance and locomotion much easier. It is only a small exaggeration to say that learning to walk is easy for humans because our legs already know how (Thelen and Smith 1993; Thelen 1995). This offloading goes beyond the boundaries of our skin. Following Gibson (1979), the natural environment is taken to be rich with affordances, or opportunities for behavior, and information that can guide behavior. As when beavers build dams, in interacting with and altering the environment, animals enhance these affordances. Kirsh and Maglio (1994) show that manipulating the environment is often an aid to problem solving. Their example is of Tetris players rotating zoids on screen, saving themselves a complicated mental rotation. Hutchins (1995) shows that social structures and well designed tools allow humans to easily accomplish tasks that would otherwise be too complex. Clark (2003) takes this further, arguing that external tools (including phones, computers, language, etc.) are so crucial to human cognition that we are literally cyborgs, partly constituted by technologies.¹

As in the 'All in the head' branch of the tree, the question of the explanatory value of mental representations arises here. Indeed, the question of the explanatory value of representations for proponents of taking cognitive systems to be combined BBEs is much more pressing, so much so that we label it 'BBE Cut 3'.

1. Wilson and Clark (2008) describe cognitive extensions like these along two (roughly orthogonal) dimensions: natural/social/technological and transient/long-term.

2.4. BBE Cut 3: Does the Brain Represent the Environment in BBE Interactions? No = Antirepresentationalism. The main explanatory purpose of positing representations in the brains of cognitive agents is to account for their connection to their environments; these internal representations are causal surrogates for distal features of the environment. But if the object of study for the cognitive sciences is the BBE system, then the environment is not distal. Hence, there is no necessary need to posit internal mental representations. Perception and cognition, on this view, are matters of direct contact between animal and environment; most intelligent behavior is real time interaction with the environment which serves “as its own model” (Brooks 1991, 139) and need not be represented. Proponents of antirepresentationalism often rely on dynamical systems theory as an explanatory tool to substantiate their claims (Beer 1995; van Gelder 1995; Chemero 2000; Thompson and Varela 2001; and see below). Proponents of antirepresentationalism are also often phenomenological realists; they argue that experience, like cognition, is a feature of BBEs rather than brains alone. Typically, they rely on continental philosophers like Heidegger and Merleau-Ponty (Varela et al. 1991) or Gibsonian ecological psychology² (Noë 2005; Chemero 2009) to make this case, though one need not.

Yes = Wide Computationalism. The majority of the proponents of cognitive science as the study of BBE interactions feel that antirepresentationalism throws out the baby with the bath water. They argue that computational cognitive science is still a valuable explanatory tool, and need not be abandoned when one spatially expands the object of study. Realizing the important role that the body and environment play in cognitive processes may alter the nature of the brain’s representations and the computations that it performs upon them, but does not change the fact that the brain is best understood as a computer. This sort of view is often called ‘wide’, ‘situated’, ‘embedded’, or ‘embodied’ computationalism (McClamrock 1995; Clark 1997; Anderson 2003; Wilson 2004). This view attempts to steer a path between GOFAI’s internalist computationalism and the more radical antirepresentationalism. It is agreed that local features of the environment sometimes need not be represented, but representations and computations still play an important explanatory role. These representations, however, are not the sentence-like representations of GOFAI; instead, they are indexical, context dependent, action oriented, and so on. Typically, they are either representations of affordances (Clark 1997) or emulations of ongoing actions (Grush 1997, 2004).

2. Indeed, most contemporary Gibsonians sit at this juncture on the wide-narrow tree, rather than as radical environmentalists.

3. Debate #2: On Explanatory Pattern. In this debate, the question is whether cognition is best explained mechanistically or dynamically. See Figure 2. It is worth noting that this debate in the cognitive and neural sciences closely mirrors one in the biological sciences (see Wilson 2004).

3.1. Is Explanation Mechanistic? YES = Mechanistic Explanations. In biological and cognitive sciences explanations frequently take the form of specification of mechanisms wherein an overall activity of a system is decomposed into subfunctions and these are localized to components of the system (Bechtel and Richardson 1993). We will focus here on the conception of mechanisms as complex systems. Machamer, Darden, and Craver define mechanisms as follows: “Mechanisms are entities and activities such that they are productive of regular change from start or setup to finish or termination conditions” (2000, 3). A mechanism is sought to explain how some process works, or how some phenomenon comes about such that the entities and activities are related to one another in a particular spatiotemporal organization. An alternative definition offered by Bechtel and Abrahamson corresponds closely to that of Machamer et al.:

A mechanism is a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism is responsible for one or more phenomena. (Bechtel and Abrahamson 2005, 423)

“To give a description of a mechanism for a phenomenon,” say Machamer et al., “is to explain that phenomenon” (2000, 3). On the complex system conception of mechanism, a mechanism is not a mere chain of objects and events, but an ontologically stable unit. Mechanisms are counterfactual-supporting and productive of determinate regularities. Each entity or part and activity or operation within a mechanism serves its role in

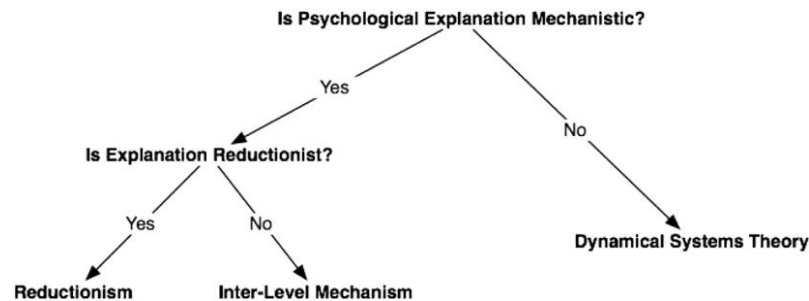


Figure 2. On explanatory style.

the bringing about of the phenomenon. That is, each part serves its function with regard to the whole. Machamer et al. stress that the activities and entities are coinstantiated within a particular mechanism—there are no activities without their entities, and no entities without their activities. Glennan (2008) stresses that mechanisms are never considered mechanisms *simpliciter*, but always mechanisms *for* a particular phenomenon. A paradigmatic example of a mechanism would be the Krebs cycle and a complete mechanistic explanation of that process would be a mechanism diagram or ‘schema’ with no black boxes left to fill in (Machamer et al. 2000).

As many have noted (Craver 2005), the connection between spatiotemporal and functional stages in a mechanism depends critically upon the spatial arrangement of the entities and upon those structured entities being oriented with respect to one another in particular ways. The temporal ordering of activities in a mechanism, as well as their rates and durations is also crucial. These spatial and temporal forms of organization gird the functional organization of mechanisms.

3.2. *Cut 1: Is Mechanistic Explanation Reductionist?* Whether or not mechanistic explanation is reductionist very much depends on what one means by reduction. This could be a question about intertheoretic reduction, the autonomy of higher-level theories or a more ontological question about the relationship between various levels, part and wholes, and so on. Here we are primarily concerned with this question: Does mechanistic explanation commit one to strong localization and decomposition? For example, can mechanisms at higher levels within the head such as large scale neural synchrony always be decomposed and localized into mechanisms at lower levels?

This is related to an issue in Debate 1, above. If it turns out that there are irreducible mechanisms (nonlocalizable or nondecomposable) at the highest levels within the brain then, because it is at roughly the same scale and there are many interactions at that shared level, it may be necessary to bring in the external environment as part of the cognitive mechanisms in question or at least as essential background for their function. As well, while we may demarcate the spatiotemporal and functional boundaries of a particular causal mechanism for various explanatory or pragmatic reasons, that same region of spacetime will often be a component in a larger mechanism or even essential background for some other mechanism. According to some, even individual brains constitute nodes in a larger sociocognitive mechanism (Deacon 1997). As Bechtel (2005) notes, mechanisms typically function only in appropriate external circumstances. Developmental biology and molecular biology have shown repeatedly, for example, that gene expression and protein production are highly context

sensitive. Wide computationalism is a good example of a mechanistic view that violates methodological individualism.

YES = Reductionism. For example, Bickle (2003), quite sanguine about the possibility of reduction in cognitive science and touting the explanatory success of cellular and molecular neuroscience, argues that psychology's memory consolidation switch that mediates the conversion of short term to long term memories reduces to the molecular mechanisms of long term potentiation (LTP). Bickle defends every variety of reductionism: the gradual and piecemeal intertheoretic reduction of cognitive psychology to neuroscience, the essential practice of methodological individualism and the eventual maximum decomposition and localization of cognitive functions to lowest-level cellular and molecular mechanisms.

NO = Interevel Mechanism. Not all agree with Bickle's characterization of the LTP case nor with his claim that such reductionism is the norm in neuroscience (Chemero and Heyser 2005; Craver 2005). Bechtel (2005) points out that experiments typically span multiple levels and require expertise in several different fields to develop a single multilevel theory. For example, even in Bickle's central case of LTP, expertise was required in molecular biology, electrophysiology, and psychology to determine the function of the NMDA receptor. Craver and Bechtel (2007) argue that mechanistic explanation guarantees higher-level theories a great deal of autonomy. They describe the situation as interfield relations that "oscillate upward and downward in a hierarchy of levels" (Craver 2005, 376). They claim that the norm in mechanistic explanation is to exhibit no predominant trend from higher toward lower levels. Multilevel descriptions of mechanisms are standard. For example,

[H]ippocampal synaptic plasticity was not discovered in the top-down search for the neural correlate of memory; rather, it was noticed in an intra-level research project that combined anatomical and electrophysiological perspectives. Such intra-level varieties of inter-field integration are far more common than their inter-level counterparts, and they are not even within the purview of reductive models of inter-field integration. Attention to the abstract structure of mechanisms and levels reveals constraints on mechanistic organization that act as loci for inter-field integration both at a given level and between levels. (Craver 2005, 374)

Craver (2005) also points out that just as the description of a mechanism is constrained by discoveries concerning higher levels, so the higher-level character of the phenomenon often must be accommodated to the findings about lower-level mechanisms. For example, we now recognize several kinds of learning and memory such as short term and long term, echoic and iconic, and explicit and implicit. Accidental or experimental damage

at the lower level in the brain tells us that these various types of memory can be decoupled from each other.

In standard accounts of reduction, the most basic level must explain all the higher-level phenomena. In mechanistic explanation however successively lower-level mechanisms account for different aspects of the phenomena in question. The idea is that ‘fields’ are integrated across levels when different fields identify various constraints on mechanisms at different levels. Different fields provide the multilevel mechanistic scaffold with a patchwork of constraints on its organization, thereby giving parameters on how the mechanism is organized. In short, “progress in understanding inter-level relations and co-evolution is more likely to be achieved if philosophers of neuroscience leave reduction behind and focus instead on the mechanistic structures that scaffold the unity of the neurosciences” (Craver 2005, 393).

Ideally, mechanisms are decomposed into component parts and operations and then each operation or function is localized in the appropriate part (Bechtel and Richardson 1993). Decomposition therefore has both a functional and structural aspect and they work hand in hand. One purpose of decomposition and localization is to include only those parts and operations that are active in the mechanism under consideration. Mechanisms have spatiotemporal and explanatory boundaries and how far these boundaries can be stretched will depend upon the mechanism under consideration. Those parts identified in the localization and decomposition of the mechanism are considered ‘inside’ and all else background. Craver and Bechtel (2007) note that the boundaries of mechanisms do not always line up with the surfaces of objects or other commonsense loci of demarcation.

One important boundary in demarcating mechanisms is the boundary of levels. Two entities or events might be at the same level if they tend to interact with each other. In this sense, entities at the same level tend to be of the same relative size, though this is not always the case. For example, when levels are defined as mechanisms functionally construed, rather than strictly structurally, larger entities will interact with smaller entities ‘at the same level’ such as cell membranes interacting with various smaller molecules. Think of the interactions between sodium ions, neural transmitters and neurons. Levels can also be further demarcated by regularities and predictabilities. The interactions of things at the same level are more predictable than the interactions of things at different levels. This heuristic claim also has obvious counterexamples. Levels are also demarcated in terms of part-whole relations. One entire mechanism can serve as a part of a ‘larger’ mechanism—one mechanism being nested within another iteratively and mechanistic explanation being recursive.

Just as the $\text{NAD}^+ \rightarrow \text{NADH}$ reactions serve as components within

the Krebs cycle, so the Krebs cycle as a whole operates as an activity within the overall mechanism of metabolism. That is, the components of the Krebs cycle are at a lower level than the Krebs cycle itself, and the Krebs cycle is in turn at a lower level than the mechanism of metabolism. Changes that occur at the higher level (in this case, metabolism) depend upon changes that occur at a lower level such as the Krebs cycle. We have also learned that the activities of a mechanism in many ways depend on what happens at higher levels (Bechtel 2005). Higher-level backgrounds essentially constrain the behavior and define the function of mechanisms.

Notice also that the part-whole relationship between mechanisms and their components is not absolute: when one thinks functionally or in terms of background conditions, x can be a component of mechanism y and y can be a component of mechanism x . For example, there is no God's-eye sense in which the Krebs cycle mechanism can be construed only as a part of the larger metabolism mechanism. It is equally valid to say that the metabolism mechanism (the sum of its various other submechanisms) is also a part of the Krebs cycle mechanism. The complexity of the demarcation between part and whole in mechanisms becomes even more twisted when we consider mechanisms that are nonlinear with feedback loops that are instantiated across various length and time scales.

This is at bottom an issue about determination. In order to adopt a strictly reductionist view about mechanisms, one would have to be convinced that the properties of the smaller components of any given mechanism (i.e., submechanisms or their parts) locally determine the capacities or properties of the mechanism as a whole. But, for example, it is not at all clear that there exists no mutual or global codetermination relations between the Krebs cycle mechanism and the metabolism mechanism as a whole, especially as these are viewed as dynamical processes with codependent evolutionary histories. Nonlinear mechanisms with system-wide feedback loops are likely to instantiate such mutual determination.

3.3. *No = Dynamical Explanation.* Although they may not have an understanding of mechanism as sophisticated as Machamer et al., Bechtel, or Craver, it seems obvious to many contemporary cognitive scientists that explanations of cognition ought to be mechanistic. Indeed, the idea that thinking is computation allows one to see how abstractions (numbers, meanings) can be encoded in a mechanical system. A growing minority of cognitive scientists, however, have eschewed mechanical explanations and embraced dynamical systems theory. That is, they have adopted the mathematical methods of nonlinear dynamical systems theory, thus employing differential equations as their primary explanatory tool. The formal apparatus of dynamical systems theory has been used to great effect to predict (within limits) the behavior of real physical systems exhibiting

various types of nonlinear effects such as so called chaotic systems. One key feature of such dynamical explanatory models is that they allow one to abstract away from causal mechanical and aggregate micro-details to predict the qualitative behavior of a class of similar systems. For example, a number of physically diverse systems exhibit the global dynamics known as the ‘period-doubling road to chaos’. The nonlinear dynamics of a cognitive system will be multiply realizable or ‘mappable’ with respect to a wide array of diverse underlying causal mechanical stories about the same processes. This is not the metaphysician’s multiple realizability of types, but the real world universality of dynamical patterns and the equations that can describe them. If models are accurate enough to describe observed phenomena and to predict what would have happened had circumstances been different, they are sufficient as explanations (van Gelder 1995; Bechtel 1998; Chemero 2000).

There are many schemes for dynamical explanation at work today in the cognitive, biological and physical sciences. By way of illustration, let us focus on coordination dynamics. Kelso and Engstrom describe coordination dynamics as

a set of context-dependent laws or rules that describe, explain and predict how patterns of coordination form, adapt, persist and change in natural systems. . . . [C]oordination dynamics seeks to identify the laws, principles, and mechanisms underlying coordinated behavior among different types of components in different kinds of systems at different levels of description. (2006, 90)

The methodology of coordination dynamics is as follows. First, for the system as a whole, discover the key coordination variables and the dynamical equations of motion that best describes how coordination patterns change over time. Second, identify the individual coordinated elements (such as neurons, organs, clapping hands, pendulums, cars, birds, bees, fish, etc.) and discern their dynamics. As Kelso and Engstrom say, this is nontrivial because the individual coordinated elements are often themselves quite complex, and are often dependent upon the larger coordinated system of which they are components (2006, 109). They put the point even more strongly, “in the complex systems of coordination dynamics, there are no purely context-independent parts from which to derive a context-independent coordinative whole” (2006, 202). Third, derive the systemic dynamics from the description of the nonlinear coupling among the elements. It is this nonlinear coupling between elements that allows one to determine connections across different levels of description. It is important to note that, as in all dynamical explanation, discovering both the systemic dynamics and that of their component parts requires specifying boundary conditions that “establish the context for particular

behaviors to arise” (Kelso and Engstrom 2006, 109). The behavior of the whole system ‘emerges’ from the nonlinear interactions among the elements of the system in a particular context where the elements and the contextual features are coupled and mutually codependent. The individual coordinating elements form a collective whole in the sense that microscopic degrees of freedom are reduced to a much smaller set of context-dependent coordination variables or order parameters that greatly constrain the behavior of the elements. Mathematical physicists often call this process ‘enslavement’ by the order parameters or the slaving principle.

Consider, for example, an experiment by Oullier et al. (2005), which uses coordination dynamics to account for spontaneous interpersonal coordination. They asked two subjects to sit across from one another, close their eyes and move their fingers up and down at a comfortable rate. Each trial was divided into three temporal segments. In one condition, subjects kept their eyes closed in segments one and three, and open in segment two. In another condition, subjects kept their eyes open in segments one and three, and closed in segment two. The same results were found in both conditions. When the subjects have their eyes closed, their finger movements were out of phase with one another. When subjects have their eyes open, their finger movements spontaneously synchronize, only to desynchronize when the subjects are asked to close their eyes again. Earlier research indicates that these finger movements are mirrored by rate-dependent activity in sensorimotor cortex (Kelso et al. 1998). So in each subject, there is a spontaneous coordination of brain activity and behavior with an external source (in this case the other subject). That is, the synchronization crosses brain, body and environment. These results are explained using coordination dynamics. In accounting for the behavior of their subjects, Oullier et al. (2005) need not worry about the mechanisms by which finger movements structure light, which impacts retinal cells, which impacts neural cells, which impacts muscles, which move fingers, and so on. The dynamical explanation of this interpersonal coordination abstracts away from all this detail, simply positing that the movements of the fingers are nonmechanically or *informationally coupled* when, and only when, the subjects have their eyes open (Kugler and Turvey 1987; Schmidt, Carello, and Turvey 1990; Kelso 1995; Kelso and Engstrom 2006). Indeed, the dynamics of the system can be modeled by attending to just one parameter—the state of the subjects’ eyes. This is the hallmark of a good dynamical explanation: the change over time of a very complex extended cognitive system, one comprising brain, body, and environment, is modeled with a comparatively simple set of equations.

4. Implications. We have just sketched the current playing field in the philosophy of cognitive science. In this section, we point to some further

connections among views and suggest some resolutions of potential conflicts. It is our hope that if the conclusions reached in this section are accepted, certain scholastic dead ends can be avoided.

4.1. Relating Positions in the Two Debates. The first thing to note is that these two debates are not completely orthogonal. That is, although your location on the wide vs. narrow decision tree does not determine your location on the explanation decision tree, there are several quite natural combinations of positions. We will review several.

4.1.1. Dynamical Explanation and Externalism. As the example from Oullier et al. (2005) indicates, many of those who prefer dynamical systems theory as an explanatory tool take cognitive systems to be wide (Kugler, Kelso, and Turvey 1980; Kugler and Turvey 1987; Kelso 1995; Port and van Gelder 1995; van Gelder 1995; Chemero 2000; Thompson 2007; Chemero and Silberstein 2008). Dynamical systems theory is especially appropriate for explaining cognition as interaction with the environment because single dynamical systems can have parameters on each side of the skin. That is, we might explain the behavior of the agent in its environment over time as coupled dynamical systems, using something like the following coupled, nonlinear equations, from Beer (1995, 1999):

$$\begin{aligned}\dot{x}_A &= A(x_A; S(x_E)), \\ \dot{x}_E &= E(x_E; M(x_A)),\end{aligned}$$

where A and E are continuous time dynamical systems, modeling the organism and its environment, respectively, and $S(x_E)$ and $M(x_A)$ are coupling functions from environmental variables to organismic parameters and from organismic variables to environmental parameters, respectively. It is only for convenience (and from habit) that we think of the organism and environment as separate; in fact, they are best thought of as comprising just one system, U . Rather than describing the way external (and internal) factors cause changes in the organism's behavior, such a model would explain the way U , the system as a whole, unfolds over time.

Of course, those who take cognitive systems to be confined to brains may also adopt dynamical explanation. Freeman and his colleagues (Skarda and Freeman 1987; Freeman and Skarda 1990) have used dynamical explanation for purely brain internal matters. In particular, they have modeled the activity in the rabbit olfactory bulb as a chaotic dynamical system whose dynamics are perturbed by odors. For Freeman, this internalism is principled: our brains are self-organizing systems whose dynamics are largely intrinsically determined. They are subject only to

very specific influences from the environment, where what these influences are is determined by the brain's own dynamics.

4.1.2. Dynamical Explanation and Representation. It is also quite common for those who prefer dynamical explanations to find mental representation unnecessary, whether they take cognition to be wide (Kugler et al. 1980; Kelso 1995; Port and van Gelder 1995; van Gelder 1995; Chemero 2000, 2009; Kelso and Engstrom 2006; Thompson 2007) or narrow (Skarda and Freeman 1987). Again, this is not a necessary connection, but is one that has been held by many theorists. Preference for dynamical explanation tends to push one away from representational explanation for two reasons. First, for fans of wide cognitive systems, it seems unnecessary to call on internal representations of environmental features when the features themselves are part of the cognitive system to be explained. Second, for internalist dynamicists like Skarda and Freeman, the dynamics of the brain are self-organizing and are perturbed only by external stimuli. Most of the dynamic structure of the brain is determined by the brain itself, so does not represent the environment. Note too that there is also pressure in the opposite direction: being an antirepresentationalist makes it difficult to embrace mechanistic explanation. This is the case because, given the toothlessness of the notion of representation as employed in the cognitive and neural sciences, nearly every mechanistic system has some parts that can be interpreted as representations. Non-mechanistic, dynamical explanations make this representation-hunting impossible. So one way to avoid representationalism is to adopt dynamical explanation. (See Bechtel 1998; Chemero 2000, 2001, 2009; Markman and Dietrich 2000; Keijzer 2001.)

4.1.3. Computationalism. Unlike dynamical systems theory and anti-representationalism, which form a semi-exclusive pair, computationalism is promiscuous. GOFAI is internalist computationalism; since many people (e.g., Churchland 1989) take connectionist networks to be kinds of computers, connectionism is also a kind of internalist computationalism; wide computationalists are externalist computationalists. This promiscuity is to be expected since computationalism was the foundational idea when cognitive science emerged as a discipline, so much so that computationalism is often called 'cognitivism'. Indeed, Wilson and Clark (2008) have suggested that computationalism is so central to the cognitive sciences that it ought to be bent to fit the data. That is, rather than admitting that cognition is not computation when cognitive systems behave like non-computers, we ought to change the definition of 'computer' so that things that had heretofore been noncomputers are captured by it. This strikes us as a very bad idea. For one thing, it threatens to make computation-

alism true by definition, when in fact it ought to be a fairly straightforwardly empirical claim. Secondly, it is out of line with the practice of contemporary cognitive science and its philosophy: the noncomputationalists in these debates (the ruthless reductionists, the large scale neural dynamicists, the radical environmentalists, the antirepresentationalists) came to their opposition to computational explanation in virtue of empirical evidence and the potential success of noncomputational methodologies.

4.2. Explanatory Pluralism. We should point out that, despite the appearances in our explanation decision tree, there is no reason that one could not be an explanatory pluralist. That is, one could believe that some cognitive phenomena are best explained reductionistically, some with an interlevel mechanistic story and some dynamically. More to the point, the very same systems can profitably be explained dynamically and mechanically. Bechtel, for example, has argued that dynamical and mechanistic approaches are complementary (1998). Furthermore, he has given a heuristic for determining which systems will resist mechanistic explanation: (1) those composed of parts that do not perform distinct operations from one another and (2) those whose behavior is chiefly the result of nonlinear interactions of the components (Bechtel and Richardson 1993; Bechtel 2002; Bechtel and Abrahamson 2005).

It is worth pausing to appreciate why nonlinearity is potentially bad news for mechanistic explanation. A dynamical system is characterized as linear or nonlinear depending on the nature of the equations of motion describing the system. A differential equation system $D(x)/Dt = Fx$ for a set of variables $x = x_1, x_2, \dots, x_n$ is linear if the matrix of coefficients F does not contain any of the variables x or functions of them; otherwise it is nonlinear. A system behaves linearly if any multiplicative change of its initial data by a factor b implies a multiplicative change of its output by b . A linear system can be decomposed into subsystems. Such decomposition fails however in the case of nonlinear systems. When the behaviors of the constituents of a system are highly coherent and correlated, the system cannot be treated even approximately as a collection of uncoupled individual parts. Instead, some particular global or nonlocal description is required, taking into account that individual constituents cannot be fully characterized without reference to larger scale structures of the system such as order parameters.

The point is that the more localizability and decomposition fail, the harder mechanistic explanation will be, and a high degree of nonlinearity is bad news for both of these. To claim that these are true of mental processes is to claim that the mind can be *decomposed* into component operations, which can then in turn be *localized* in brain areas. It is still

an open question whether or not the brain is so radically nonlinear as to defeat mechanistic explanation. Cosmelli et al. (2007) reviews data on large scale neural dynamics and synchrony that suggests that localization and decomposability, hence mechanical explanation, will fail. Bechtel, on the other hand, thinks there is cause for hope because “the strategy of decomposition and localization is compatible with discovering great amounts of interactivity that gives rise to complex dynamics” (2002, 239). More importantly, Bechtel points out that dynamical and mechanistic explanations are mutually supporting in that information about the mechanisms and their components help us to better formulate the dynamics and insight into the dynamics can help us discover key mechanisms (2002, 240). Similarly, Clark (1997) has argued that dynamical and computational (i.e., mechanistic) explanatory patterns ought to ‘interlock’ in a complete explanation of cognition. Even dynamicists Kelso and Engstrom point out that dynamic patterns require pattern generators, that is, causal mechanisms, and vice versa (2006, 95). There is much to be said for pluralism about explanation. Animal behavior and animal brains are very complex, and we can see no *a priori* reason that all aspects of them or any one aspect of them, ought to be explained in *any* one way, whether or not explanations interlock or are complementary. For example, it is important to pursue both modular and relatively nonmodular explanations of brain function, because all the evidence to date is that brain function is *both* localizable and highly integrated.

Given the discussion above, this quite sensible explanatory pluralism might be unavailable to the antirepresentationalist, who must eschew mechanistic explanation to avoid representation hunting. Antirepresentationalists, like everyone else, should be explanatory pluralists. This would indicate that they need some other tool to argue with representationalists like Bechtel (1998) and Markman and Dietrich (2000), who find representations in their purportedly representation-free models of cognition. As far as we know, no such tool is presently available. In fact, antirepresentationalist cognitive scientists have begun to ignore representation hunters, insisting that the representation hunters’ understanding of representation is empty (Wheeler 2005a). In so doing, they are refusing, admirably in our view, to participate in a scholastic debate.

On our view, dynamical and mechanistic explanation of the same complex system get at different but related features of said system described at different levels of abstraction and with different questions in mind. We see no *a priori* reason to claim that either kind of explanation is more fundamental than the other. In fact, because both types of explanation have become so much more holistic or wide in certain cases, we think they share several features in common that give us hope for a unified science of complex systems. For example, as with coordination dynamics,

we have seen that modern mechanistic explanation places much more emphasis on the dynamic organization, contextual features and interrelations of a given mechanism.

Take self-organized natural systems as an example. Whether viewed mechanistically or dynamically, thermodynamically open, far from equilibrium, nonlinear complex systems such as those found in neural and behavioral dynamics can be characterized as self-organizing systems generated by complex feedback mechanisms and relationships. The key features of these self-organizing systems are as follows. First, they arise without homunculus-like agents assembling them, and we need posit no new laws of nature to explain them. Self-organizing systems are thus spontaneously arising coordination patterns. Second, such self-organizing patterns are generated by context-dependent order parameters. Third, self-organizing systems result from both short range and long range informational, causal-mechanical, functional, and other kinds of coupling. This coupling can occur among elements within a part of the system (e.g., among neurons in a particular brain area), between parts of the same system (e.g., between brain areas), and between various features of system and its environment (Kelso and Engstrom 2006, 91–93, 112).

Fourth, self-organizing systems maintain autonomy, boundaries and identity via certain self-maintenance operations of some sort, what some describe as autopoiesis (Varela 1979). This autonomy is a result of the fact that self-organizing systems constrain the behavior of their parts such that their organization yields a functional whole even across many internal changes and varying environmental stimuli. As Varela puts it: “Autopoietic machines are autonomous: that is, they subordinate all changes to the maintenance of their own organization, independently of how profoundly they may be otherwise transformed in the process” (1979, 15). Varela also provides a more generalized definition of autonomous systems:

Autonomous systems are mechanistic (dynamical) systems defined as a unity by their organization. We shall say that autonomous systems are organizationally closed. That is, their organization is characterized by processes such that (1) the processes are related as a network, so that they recursively depend on each other in the generation and realization of the processes themselves, and (2) they constitute the system as a unity recognizable in the space (domain) in which the processes exist. (1979, 55)

The more intrinsically integrated and the more nonlinear a self-organizing system is, the more localizability and decomposability will fail. Simple systems are decomposable; that is, their components are not altered by construction, and are recoverable by disassembly. It is not necessarily so for self-organizing systems. Once a system becomes truly autopoietic,

component systems lose their autonomy. As well, because reproduction of components is now dependent on the new intrinsically integrated whole, the generation of mechanisms to prevent component autonomy is favored, as in the case of multicellularity.

Regardless of how we characterize them, dynamically or mechanistically, the relative failure of localizability and decomposability in self-organizing systems implies that the deepest explanation for such systems cannot be in terms of the Lego-philosophy of atomism or mechanism. That is, both dynamical and mechanistic explanations of self-organized systems will conflict with intertheoretic and mereological reductionism, at least as these are typically conceived by philosophers. In order to get mereological type-reductionism off the ground, there have to be context-independent or invariant fundamental parts with intrinsic properties such as atoms or cells whose temporal evolution is governed by some fundamental context-independent dynamical laws. On this reductionist view, all the other ‘higher-level’ features of the atomistic system, such as the coordination dynamics it exhibits will be determined by the fundamental constituents and the laws governing them—the higher-level features will come for free given the fundamental facts. If one views nonlinear dynamical systems as at bottom structured aggregates of physical parts then one must agree with Kellert:

It is important to clarify that chaos theory argues against the universal applicability of the method of micro-reductionism [such as deriving coordination dynamics from say Newton’s laws], but not against the validity of the philosophical doctrine of reductionism. That doctrine states that all properties of a system are reducible to the properties of its parts. . . . Chaos theory gives no examples of ‘holistic’ properties which could serve as counter-examples to such a claim. (1993, 90)

However, as we have seen, self-organizing systems—no matter how characterized—are not structured aggregates of physical parts (Silberstein 1999). Dynamical coordination patterns emerge from the collective behavior of coupled elements in a particular context and in turn the behavior of each individual element is constrained by the collective behavior of the whole in that context. Kelso and Engstrom say that

in the complex systems of coordination dynamics, there are no purely context-independent parts from which to derive a context-independent coordination whole, even though we often try and occasionally succeed to analyze them as such. Again, in coordination dynamics, as in the brain itself, the parts are not one nor are they separate. (2006, 202)

The kind of explanatory pluralism and ontological holism found in self-organizing systems is sure to be bad news for vitalism, dualism, and preformationism of any kind; however, it is also bad news for mereological reductionism. And as we saw earlier, Bechtel makes it clear that actual mechanistic explanatory patterns and methods are bad news for philosophical accounts of intertheoretic reduction such as the D-N model.

4.3. Causal Pluralism. As a corollary to explanatory pluralism, we also endorse causal pluralism. Proponents of strictly mechanistic explanation often suggest that dynamical explanations are mere descriptions of phenomena. They say *what* happens, but not *why* it happens or what *causes* it to happen. The implication here is that possessing the deepest explanation of a phenomenon requires filling in the dynamical explanation with a mechanistic explanation. This may be true if we are limited to efficient causation, that is, if we assume that efficient causation is the most fundamental kind of causal explanation. But proponents of dynamical explanation also often make use of formal causation. In a typical self-organizing dynamical system, the emergent features of the system arise from the behavior of the coupled elements, yet also constrain their behavior. The simplest, and best understood system we can use to make this case is Rayleigh-Bénard convection (Bishop 2008). Rest assured that the relevant features of this case do map onto dynamical systems explanations in cognitive science and neuroscience.

Rayleigh-Bénard convection occurs in a fluid such as oil between two plates differing in temperature. When the difference in temperature, Δt , is large enough to produce convection, there is a breakup of the stable conductive state and large scale rotating structures resembling a series of parallel cylinders called Bénard cells are eventually produced. Associated with this new stable pattern is a large scale, nonlocal constraint on the individual motions of fluid elements due to a balancing among effects owing to the structural relations of each fluid element to all other fluid elements, system boundaries, and so on. If Δt is further increased beyond a certain threshold, the cells begin to oscillate transversely in complicated ways. Further increases in Δt lead to destruction of the cells and to chaotic behavior in the fluid, where spatial correlations persist. However, when Δt is sufficiently large, fluid motion becomes turbulent and uncorrelated at any two points in the fluid. Rayleigh-Bénard convection under increasing Δt exhibits typical phenomena associated with chaotic nonlinear systems, such as period-doubling cascades, phase locking between distinct oscillatory modes, and sensitive dependence on initial conditions.

Fluid elements participate in particular Bénard cells within the confines of the container walls. The system as a whole displays integrity as the constituents of various hierarchic levels exhibit highly coordinated, co-

hesive behavior. Additionally, the organizational unity of the system is stable to small perturbations in temperature and adapts to larger changes within a particular range of differences in temperature. Furthermore, Bénard cells act as a control hierarchy, constraining the motion of fluid elements. Bénard cells emerge out of the motion of fluid elements as Δt exceeds the appropriate threshold, but these large scale structures determine modifications of the configurational degrees of freedom of fluid elements such that some motions possible in the equilibrium state are no longer available. In the new nonequilibrium steady state, the fluid elements exhibit coherent motion (Bénard cells), but most of the states of motion characteristic of the original uniform state are no longer accessible. As Bishop puts it,

[I]t can be shown that the uniform and damped modes follow—or are enslaved by—the unstable growing mode leading to cell formation. Basically, the variables characterizing steady and damped modes can be systematically replaced by variables characterizing the growing mode; consequently there is a corresponding reduction in the degrees of freedom characterizing the system's behavior. Such a reduction of the degrees of freedom is crucial to the large-scale behavior of convection systems. (2008, 238)

How should we view the causal relations between fluid elements and Bénard cells? In complex systems, levels of structure are often only distinguishable in terms of dynamical time scales and are coupled to each other in such a way that at least some of the higher-level structures are not fully determined by, and even influence and constrain the behavior of, constituents in lower-level structures. The fluid elements are of course necessary for the existence of Bénard cells; however, the dynamics of fluid elements themselves are not sufficient to bring about the cell structures—context plays a central role here. Each fluid element is coupled with all others, reflecting large scale structural constraints on the dynamics of the fluid elements through such relations as enslavement and bulk flow. Surprisingly, Bénard cells even modulate their fluid element and energy intake in order to maintain cohesion within a range of variations in differences in temperature.

At least during the process of pattern formation, the relationship between fluid elements and cells is not a simple matter of a feedback mechanism. A feedback mechanism typically involves a reference state and a number of feedback-regulated set points, so that comparisons between the evolving and reference states can be made and adjustments effected—consider a thermostat. But during pattern formation in Rayleigh-Bénard convection, the appropriate reference states do not exist. The original

uniform state is destroyed and a new nonequilibrium steady state forms the Bénard cells.

Bénard cells arise from the dynamics of fluid elements in the Rayleigh-Bénard system as a function of Δt (a key order parameter), such that each fluid element becomes coupled with every other fluid element. Bénard cells act as a control hierarchy, constraining and modifying the trajectories of fluid elements. So although the fluid elements are necessary to the existence and dynamics of Bénard cells, they are not sufficient to determine their dynamics, nor are the fluid elements sufficient to fully determine their own motions. Rather, the large scale structure supplies a governing influence constraining the local dynamics of the fluid elements. Bénard cells are the formal causes of the altered dynamics of the fluid elements and the fluid elements are one of the efficient causes of the cells. The properties of integrity, integration, and stability exhibited by Bénard cells are determined by the dynamic properties of the nonlocal relations of all fluid elements to each other. This kind of formal causation by emergent features is a routine property of dynamical models: emergent, formal causes are modeled as collective, control variables. Furthermore, it is apparently ubiquitous in nature and wholly unmysterious. If simple systems possessing no biological or cognitive functions such as Rayleigh-Bénard convection exhibit formal causes then there is every reason to believe that both brains (Bressler and Kelso 2001; Thompson and Varela 2001) and BBE systems (Kelso 1995) have emergent features that act as even more sophisticated formal causes. So we see no reason not to embrace causal pluralism as a corollary to explanatory pluralism. And this takes nothing away from the fact that Rayleigh-Bénard convection can be and is treated dynamically and mechanistically (see Bishop 2008 for details on both types of explanation; see also Silberstein 2002).

4.4. Toward a Holistic Cognitive Science. We have seen that when it comes to complex, self-organizing systems, the dynamical and mechanistic perspectives seem to be somewhat convergent. These types of explanation complement one another. We often devise dynamical explanations of a system when we cannot work out all the mechanisms and we often pursue mechanistic explanations when the number of variables or the nonlinearity of a system makes dynamical predictions intractable. Moreover, both of these perspectives are at odds with philosophical models of reductionism. This parallels the debate over methodological individualism: taking the object that the cognitive and neural sciences ought to study to be the BBE system is also at odds with reductionism.

There are of course many gradations of both positions, ultimately shading off into one another. Individualists can be more or less holistic, for example. Even having decided that good cognitive and neuroscience must

confine itself to the boundaries of the head, there still remains the question of which scale of cognitive or brain activity to pitch the explanation at. At what 'level' should we explain cognitive systems? Those explanations involving the more basic elements of a system (such as a single neuron) and the purportedly intrinsic or local properties of those elements are the most deeply individualistic in kind. Individualist explanations that focus on large scale and inherently relational features of cognitive systems such as functional features or large scale neural dynamics are the least individualistic. For example, Bickle (2003) and Skarda and Freeman (1987) embrace methodological individualism. But the more thoroughly individualistic Bickle focuses on reductionist explanation and nonrelational properties at the microscale, while Skarda and Freeman focus on dynamical explanations of relational and systemic properties. However it is an interesting question just how far holistic or externalist science can get.

While it is not easy to define holistic science, let us try to capture it in two dimensions. First, we can say that holistic science is relatively wide, externalist and methodologically holistic in the ways we have discussed. BBE cognitive science is a good example. Second, we can say that holistic science abstracts relatively far away from the actual micro-details of the system in question. Attempts to construct broadly applicable dynamical systems theory explanations, such as Kelso's coordination dynamics that apply equally to ants, cars, neurons, genes, many-bodied quantum and classical systems, galaxy clusters, and so on, is a good example of this kind of science. How far can such science get? Perhaps we can say that the more broadly applicable a scientific explanation or theory is and the more localization and decomposition fail in its explanatory patterns, the more holistic it is.

The biggest advantage of methodological individualism and methodological reductionism is that they have proven to be phenomenally successful tools for prediction, explanation, and intervention in physical, chemical and biological systems. As Lewontin puts it,

This [methodological reductionism] was a radical departure from the holistic pre-Enlightenment view of natural systems as indissoluble wholes that could not be understood by being taken apart into bits and pieces. . . . Over the last three hundred years the analytic model has been immensely successful in explaining nature in such a way as to allow us to manipulate and predict it. (2000, 72)

Indeed, it is probably fair to say that the only reason more holistic scientific methodologies are being developed, be they dynamical or mechanistic, is that methodological reductionism has hit the wall in many scientific disciplines involving the study of complex, nonlinear, multileveled, nested and self-organizing systems. In fact, as we have seen, even the most de-

tailed causal mechanical explanations of such systems are much more holistic than is suggested by methodological individualism. Causal mechanical explanations have historically been viewed as the contrast case with holistic explanation. One point of the previous discussion of modern day causal mechanical explanation is that the nature of complex cognitive and biological systems is such that sophisticated causal mechanical explanations of them are holistic in a number of ways, threatening methodological individualism even in the realm of mechanism. This ought to make it easier to appreciate why there is renewed interest in developing rigorous holistic science.

However, the biggest pragmatic or practical problem with developing holistic science is obvious: explanatory and predictive successes are hard to come by when dealing with complex systems. The principle worry here is that too much holism makes science impossible. For given enough holism, what is thought to be a 'system' or 'subsystem' is, after all, just a convention. Science seems to be predicated on dividing the world objectively into the relevant explanatory parts or systems per the explanatory task at hand: "obscurantist holism is both fruitless and wrong as a description of the world" (Lewontin 2000, 111). One of the major things that separates science from mysticism such as Hindu nondualism or pseudoscience such as astrology, is the twin constraints of methodological and metaphysical individualism. So why pursue holistic science at all? Lewontin answers:

It seems abundantly clear to us now that the [obscurantist] holistic view of the world obstructs any possibility of a practical understanding of natural phenomena. But the success of the clock model, in contrast to the failure of obscurantist holism, has led to an overly simplified view of the relations of parts to wholes and of causes to effects. Part of the success of naïve reductionism and simplistic analysis comes from the opportunistic nature of scientific work. Scientists pursue precisely those problems that yield to their methods, like a medieval army that besieges cities for a period, subduing those whose defenses are weak, but leaving behind, still unconquered, islands of resistance. . . . Successful scientists soon learn to pose only those problems that are likely to be solved. Pointing to their undoubted successes in dealing with the relatively easy problems, they assure us that eventually the same methods will triumph over the harder ones. (2000, 72–73)

Lewontin's point is that methodological and metaphysical individualism constitute a kind of circular reasoning or question begging perspective unless we test them by pursuing the possibility of more holistic science.

Our recommendation, then, for doing cognitive science and the phi-

philosophy of cognitive science is that it be done with the constraints of explanatory and causal pluralism. In the long run, we must judge the relative success of reductionist and holistic science (and everything in between) by their explanatory successes and failures. If we were to call the race now, the winner would probably be on the reductionist end of the spectrum. However, it is much too early to make that call. It is also probably clear to the reader by now that we believe that holistic science has had enough success to make it worth pursuing in the future.

5. Conclusion. We hope to have shown that the successor to the philosophy of mind promises to be extremely rich. Indeed, there are many remaining issues in these debates that we have not touched upon, such as the explanation and status of consciousness and phenomenology. Quine's dictum, that philosophy of science is philosophy enough, leaves plenty of work for philosophers of cognitive science, and saves them from worrying about the c-fibers of Martian zombies or the beliefs of SwampMan.

REFERENCES

- Anderson, M. (2003), "Embodied Cognition: A Field Guide", *Artificial Intelligence* 149: 91–130.
- Bechtel, W. (1998), "Representations and Cognitive Explanations: Assessing the Dynamist's Challenge in Cognitive Science", *Cognitive Science* 22: 295–318.
- (2002), "Decomposing the Brain: A Long Term Pursuit", *Brain and Mind* 3: 229–242.
- (2005), "Mental Mechanisms: What Are the Operations?", in B. Bara, L. Barsalou, and M. Bucciarelli (eds.), *Proceedings of the 27th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum, 208–213.
- Bechtel, W., and A. Abrahamson (2005), "Explanation: A Mechanist Alternative", *Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 421–444.
- Bechtel, W., and R. C. Richardson (1993), *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton, NJ: Princeton University Press.
- Beer, R. (1995), "Computational and Dynamical Languages for Autonomous Agents", in R. F. Port and T. van Gelder (eds.), *Mind as Motion*. Cambridge, MA: MIT Press, 121–147.
- (1999), "Dynamical Approaches to Cognitive Science", *Trends in Cognitive Sciences* 4: 91–99.
- Bickle, J. (2003), *Philosophy and Neuroscience: A Ruthlessly Reductive Account*. Dordrecht: Kluwer Academic.
- Bishop, R. (2008), "Downward Causation in Fluid Convection", *Synthese* 160: 229–248.
- Bressler, S. L., and J. A. S. Kelso (2001), "Cortical Coordination Dynamics and Cognition", *Trends in Cognitive Sciences* 5: 26–36.
- Brooks, R. (1991), "Intelligence without Representation", *Artificial Intelligence* 47: 139–159.
- Chemero, A. (2000), "Anti-representationalism and the Dynamical Stance", *Philosophy of Science* 67: 625–647.
- (2001), "Dynamical Explanation and Mental Representation", *Trends in Cognitive Sciences* 5: 140–141.
- (2009), *Radical Embodied Cognitive Science*. Cambridge, MA: MIT Press, forthcoming.

- Chemero, A., and C. Heyser (2005), "Object Exploration and the Problem with Reductionism", *Synthese* 147: 403–423.
- Chemero, A., and M. Silberstein (2008), "Defending Extended Cognition," in V. Sloutsky, B. Love, and K. McRae (eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society*. New York: Psychology Press.
- Churchland, P. (1989), *A Neurocomputational Perspective*. Cambridge, MA: MIT Press.
- Clark, A. (1997), *Being There*. Cambridge, MA: MIT Press.
- (2003), *Natural Born Cyborgs*. New York: Oxford University Press.
- Cosmelli, D., J.-P. Lachaux, and E. Thompson (2007), "Neurodynamics of Consciousness", in P. D. Zelazo, M. Moscovitch, and E. Thompson, (eds.), *The Cambridge Handbook of Consciousness*. Cambridge: Cambridge University Press, 730–770.
- Craver, C. (2005), "Beyond Reduction: Mechanisms, Multifid Integration and the Unity of Neuroscience", *Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 373–395.
- Craver, C. F., and W. Bechtel (2007), "Top-Down Causation without Top-Down Causes", *Biology and Philosophy* 22: 547–563.
- Deacon, T. (1997), *The Symbolic Species: The Co-evolution of Language and the Brain*. New York: Norton.
- Fodor, J. (1975), *The Language of Thought*. Cambridge, MA: Harvard University Press.
- Fodor, J., and Z. Pylyshyn (1988), "Connectionism and the Cognitive Architecture", *Cognition* 28: 3–71.
- Freeman, W. (1999), *How Brains Make Up Their Minds*. New York: Columbia University Press.
- Freeman, W., and C. Skarda (1990), "Representations: Who Needs Them?," in J. L. McGaugh et al. (eds.), *Brain Organization and Memory Cells, Systems and Circuits*. New York: Oxford University Press, 375–380.
- Gibson, J. (1979), *The Ecological Approach to Visual Perception*. Hillsdale, NJ: Erlbaum.
- Glennan, S. (2008), "Mechanisms", in S. Psillos and M. Curd (eds.), *The Routledge Companion to the Philosophy of Science*, forthcoming.
- Grush, R. (1997), "The Architecture of Representation", *Philosophical Psychology* 10: 5–24.
- (2004), "The Emulation Theory of Representation: Motor Control, Imagery, and Perception", *Behavioral and Brain Sciences* 27: 377–442.
- Haugeland, J. (1985), *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press.
- Hineline, P. (2006), "Multiple Scales of Process, and the Principle of Adduction", in E. Ribes-Inesta and J. E. Burgos (eds.), *Knowledge, Cognition, and Behavior*. Guadalajara: University of Guadalajara Press, Chapter 11.
- Hutchins, E. (1995), *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Keijzer, F. (2001), *Representation and Behavior*. Cambridge, MA: MIT Press.
- Kellert, S. (1993), *In the Wake of Chaos*. Chicago: University of Chicago Press.
- Kelso, J. (1995), *Dynamic Patterns: The Self-Organization of Brain and Behavior*. Cambridge, MA: MIT Press.
- Kelso, J., and D. Engstrom (2006), *The Complementary Nature*. Cambridge, MA: MIT Press.
- Kelso, J. A. S., A. Fuchs, T. Holroyd, R. Lancaster, D. Cheyne, and H. Weinberg (1998), "Dynamic Cortical Activity in the Human Brain Reveals Motor Equivalence", *Nature* 23: 814–818.
- Kirsh, D., and P. Maglio (1994), "On Distinguishing Epistemic from Pragmatic Action", *Cognitive Science* 18: 513–549.
- Kugler, P. N., J. A. S. Kelso, and M. T. Turvey (1980), "Coordinative Structures as Dissipative Structures: I. Theoretical Lines of Convergence", in G.E. Stelmach and J. Requin (eds.), *Tutorials in Motor Behavior*. Amsterdam: North-Holland, 3–47.
- Kugler, P., and M. Turvey (1987), *Information, Natural Law, and the Self-Assembly of Rhythmic Movement*. Hillsdale, NJ: Erlbaum.
- Lewontin, R. (2000), *The Triple Helix: Gene, Organism and Environment*. Cambridge, MA: Harvard University Press.

- Machamer, P., L. Darden, and C. Craver (2000), "Thinking about Mechanisms", *Philosophy of Science* 67: 1–25.
- Markman, A., and E. Dietrich (2000), "In Defense of Representation", *Trends in Cognitive Sciences* 4: 470–475.
- McClamrock, R. (1995), *Existential Cognition*. Chicago: University of Chicago Press.
- Noë, A. (2005), *Action in Perception*. Cambridge, MA: MIT Press.
- Oullier, O., et al. (2005), "Spontaneous Interpersonal Synchronization", in C. Peham, W. I. Schöllhorn, and W. Verwey (eds.), *European Workshop on Movement Sciences: Mechanics-Physiology-Psychology*. Cologne: Sportverlag, 34–35.
- Port, R., and T. van Gelder (1995), *Mind as Motion*. Cambridge, MA: MIT Press.
- Quine, W. V. O. (1966), *The Ways of Paradox and Other Essays*. New York: Random House.
- Ramsey, W. (1997), "Do Connectionist Representations Earn Their Explanatory Keep?", *Mind and Language* 12: 34–66.
- Rumelhart, D., et al. (1986), "Schemata and Sequential Thought Processes in PDP Models", in D. Rumelhart, J. McClelland, and the PDP Research Group (eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 2, *Psychological and Biological Models*. Cambridge, MA: MIT Press, 7–57.
- Schmidt, R. C., C. Carello, and M. T. Turvey (1990), "Phase Transitions and Critical Fluctuations in the Visual Coordination of Rhythmic Movements between People", *Journal of Experimental Psychology: Human Perception and Performance* 16: 227–247.
- Silberstein, M. (1999), "The Search for Ontological Emergence", *Philosophical Quarterly* 49: 182–200.
- (2002), "Reduction, Emergence, and Explanation", in P. Machamer and M. Silberstein (eds.), *The Blackwell Guide to the Philosophy of Science*. Oxford: Blackwell, 80–107.
- Skarda, C., and W. Freeman (1987), "How the Brain Makes Chaos to Make Sense of the World", *Behavioral and Brain Sciences* 10: 161–195.
- Thelen, E. (1995), "Time-Scale Dynamics and the Development of Embodied Cognition", in R. F. Port and Timothy van Gelder (eds.), *Mind as Motion: Exploration in the Dynamics of Cognition*. Cambridge, MA: MIT Press, 69–100.
- Thelen, E., and L. Smith (1993), *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press.
- Thompson, E. (2007), *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Cambridge, MA: Harvard University Press.
- Thompson, E., and F. Varela. (2001), "Radical Embodiment: Neural Dynamics and Consciousness", *Trends in Cognitive Sciences* 5: 418–425.
- van Gelder, T. (1995), "What Might Cognition Be If Not Computation?", *Journal of Philosophy* 91: 345–381.
- Varela, F. J. (1979), *Principles of Biological Autonomy*. New York: Elsevier North-Holland.
- Varela, F., E. Thompson, and E. Rosch (1991), *The Embodied Mind*. Cambridge, MA: MIT Press.
- Wheeler, M. (2005a), "Friends Reunited? Evolutionary Robotics and Representational Explanation", *Artificial Life* 11: 215–232.
- (2005b), *Reconstructing the Cognitive World*. Cambridge, MA: MIT Press.
- Wilson, R. (2004), *Boundaries of the Mind*. New York: Cambridge University Press.
- Wilson, R., and A. Clark (2008), "How to Situate Cognition: Letting Nature Take Its Course", in M. Aydede and P. Robbins (eds.), *The Cambridge Handbook of Situated Cognition*. Cambridge: Cambridge University Press, forthcoming.