

Structure and Application of Dynamical Models in Cognitive Science

Maurice Lamb (Lambmi@mail.uc.edu)

Department of Philosophy, 206 McMicken Hall, University of Cincinnati
Cincinnati, OH 45221 USA

Anthony Chemero (Anthony.chemero@uc.edu)

Department of Philosophy and Department of Psychology, 206 McMicken Hall, University of Cincinnati
Cincinnati, OH 45221 USA

Abstract

In philosophy of science, Neo-mechanists argue that explanations are only successful when formulated in terms of the behaviors of discrete decomposable components that constitute the system of interest. This approach to explanation implicitly denies the significance of non-linear interactions in structuring the behavior of complex cognitive systems. Recently, Neo-mechanists have claimed that JAS Kelso and colleagues have begun to favor neo-mechanistic explanations of neuroscientific phenomena; particularly in the application of the neural field model to rhythmic coordination behaviors. We will argue that this view is the result of a failure to understand dynamic systems explanations and the general structure of dynamic systems research. Further, we argue that the explanations cited are in fact not neo-mechanistic explanations. In this paper, we will show that these neo-mechanists have misunderstood the work by Kelso and colleagues, which blunts the force of one of their arguments.

Keywords: Explanation; Dynamic systems; Mechanism; HKB; Neural Field Model; Tripartite Scheme; Neuroscience

Introduction

Many scientists have some criteria for deeming some findings as explanatory and others as useful but not explanatory, though these criteria are rarely formalized. Attempts at defining a simple account of explanation in terms of necessary and sufficient conditions have often come up short. In the philosophy of science, a theory of explanation, referred to as the neo-mechanist approach, has been developed in terms of a particular understanding of mechanisms in scientific investigations. In the context of cognitive science, Neo-mechanists (Bechtel and Abrahamsen 2005; Craver 2007; Bechtel 2009, 2011; Kaplan and Bechtel 2011; Kaplan and Craver 2011) argue that in order for a claim to be an explanation in cognitive science it must reveal something about the decomposable mechanisms of a cognitive system. As part of their arguments, they claim that JAS Kelso and colleagues working on cognitive systems are shifting away from dynamic systems explanations of cognitive and behavioral phenomena in favor of neo-mechanistic explanation of neuroscientific phenomena (Kaplan and Bechtel, 2011; Kaplan and Craver 2011). We will argue that this view is the result of a failure to understand dynamic systems explanations and the general structure of dynamic systems research, and that the explanations by Kelso and colleagues cited are in fact not neo-mechanistic explanations. We will

not, in this very short paper, offer a full defense of dynamical models as genuine explanations. Here, we will simply show that these neo-mechanists have misunderstood the work by Kelso and colleagues, which blunts the force of one of their arguments.

Our paper will proceed in three parts. First, we will briefly describe neo-mechanisms and what we will call “The Scott Freaking Kelso Argument”. Second, we will outline the basic methodology of one form of dynamic systems research. In this section our aim is to clarify the structure and formulation of dynamic systems models in the context of Synergetics in order to distinguish this strand of dynamic systems research from neo-mechanistic theories. Third, we will examine the role of the neural field model in dynamic systems research (Jirsa and Haken, 1996; Jirsa et al. 1998; Jantzen et al. 2009). This work has been cited as a supposedly clear example that dynamic systems researchers ultimately depend on neo-mechanistic explanations to make their models explanatory (Kaplan and Bechtel 2011; Kaplan and Craver 2011). We will show that the neural field model is a dynamic systems model, and thus, application of the neural field model is continuous with dynamic systems theory not contrary to it.

Neo-Mechanism and the SFK Argument

For the past 20 years, there has been a consensus among philosophers of science that mechanistic explanation is important in the life sciences. Bechtel and Richardson (2010) defined neo-mechanistic explanation as explanation that involves decomposing some phenomenon into component operations, and then localizing those component operations in physiological structures of organisms. These component operations are taken to produce or to be responsible or to account for the phenomena. Decomposition involves developing a model of a system’s behavior by identifying discrete component parts and their linear, or weakly non-linear, interactions. While it cannot be argued here, non-linearity in a system is not trivial. Genuinely non-mechanistic descriptive and explanatory strategies are required to capture features of non-linear interactions that are otherwise unavailable to neo-mechanistic paradigms of explanation. Many of these non-linear features are at the center of debates over emergence in natural systems, though these debates are outside the scope of the current project. Localization involves mapping those discrete components and interactions onto features of a

physical system (Bechtel and Richardson, 2010). It is important to recognize that the just the application of any type of model to a physical system is not a neo-mechanistic explanation. The model must be decomposable, i.e. consist of discrete components and primarily linear couplings.¹ We acknowledge that neo-mechanistic approaches to physical systems provide an important style of explanation in biology, neuroscience, psychology, and cognitive science. However, neo-mechanistic explanation will be ineffective for systems that are not decomposable or nearly decomposable, due to, for example, high degrees of nonlinearity and thus must be augmented by other explanatory approaches (Bechtel and Richardson, 2010). For this reason, we think that both approaches are useful in distinct contexts. Recently, however, proponents of neo-mechanistic explanation have argued that explanations in the life science just are neo-mechanistic explanations. This has been argued in one of two ways: either it is argued that non-mechanistic explanations are not really explanations, or it is argued that apparently non-mechanistic explanations are in fact neo-mechanistic explanations. Many neo-mechanists make both arguments (e.g., Kaplan and Craver 2011, Kaplan and Bechtel 2011, Bechtel 2011, Bechtel and Abrahamsen 2010).

As noted above, here we will address one argument that non-mechanistic explanations are not genuine explanations. (We address other arguments elsewhere.) Here is a quote from Kaplan and Craver 2011,

Many proponents of dynamic systems theory such as Kelso now appear to recognize the importance of mechanistic explanation. After developing the HKB model, Kelso and colleagues began researching how this behavioral regularity results from features of the underlying organization of component neural systems and their dynamics (see, e.g., Schöner and Kelso 1988; Jirsa et al. 1998; Jantzen et al. 2009). Kelso and colleagues (Jirsa et al. 1998) recently proposed a *neural field* model connecting the observed phase shift described by HKB to the underlying dynamics of neural populations in motor cortex.

We call this the “Scott Freaking Kelso Argument” (SFK). It also appears in Kaplan and Bechtel (2011). The basic rhetorical move is that dynamical models are so obviously not genuine explanations the way neo-mechanistic models are that even *Scott Freaking Kelso* is pursuing neo-mechanistic explanations. Compare an equally false claim: “Many evolutionary biologist such as Richard Dawkins now appear to recognize Jesus Christ as their personal savior.” Although, we would never recommend this to introductory

¹ We believe some confusion enters philosophical debates of explanation because many scientists use the term ‘mechanism’ to refer to the claim that features of the model map onto features of a physical system. We do not take issue with this usage, since common usage does not pre-theoretically eliminate the possibility that macro-level patterns of system behavior could constrain micro-level patterns of behavior.

logic students as a valid argument, its rhetorical force is undeniable. If even the pioneers of dynamical modeling in the cognitive sciences—the people who understand them better than anyone else—don’t believe that dynamical models are genuine explanations, we probably shouldn’t either. Crucially, the SFK argument depends upon whether the neural field model (Jirsa et al 1998) really is an attempt to “transform a merely descriptive model into a mechanistic one”. In the remainder of this paper, we argue that it is not.

Dynamic Systems and the Tripartite Scheme

Dynamic systems theory may be used to refer a family of research orientations that model the behaviors of non-linear systems over time using differential equations. Our focus in this paper is on Synergetics, as exemplified in the widely discussed Haken-Kelso-Bunz (HKB) model and structured by what Kelso calls the Tripartite Scheme (Kelso, 1995; Haken et al, 1985). The Tripartite Scheme involves a three part approach to dynamic systems models (See Figure 1). The three parts are the boundary constraints, the collective variable, and the coordination components.² Each part of the Tripartite Scheme provides a distinct representation of the system and behavior of interest but, crucially, no part is independent of the others.

The *boundary constraints* define the system and behavior to be observed by defining the initial conditions of system’s behavior and defining the relevant variables and parameters. In the HKB model as applied to rhythmic finger coordination, the initial conditions are defined as the oscillations of an individual’s left and right index fingers at a particular frequency. The relevant variables and parameters include changes in oscillation frequency, changes in amplitude, spatial orientation, etc... The boundary constraints define the range in which the system of interest is observed. Changing the boundary constraints, changes the system of interest. Generally, dynamic systems models are sensitive to changes in both initial and parametric values. For example, given a specified parameterization the initial conditions may determine whether the system exhibits chaotic or non-chaotic behavior. Likewise, changes in the parameters can result in different bifurcations as well as distinct patterns of behavior. Thus, the same limbs or cluster of neurons may have certain coordination dynamics under one set of boundary constraints and different coordination dynamics under a different set of boundary constraints.

The *collective variables* characterize the coordination dynamics of the system and its behavior. The collective variables represent a relationship(s) between the coordination components. The representation of the system

² There is some variation in the terms Kelso and colleagues use across publications to describe these parts: boundary constraints are sometimes referred to as task constraints and the collective variable is sometimes referred to as the order parameter.

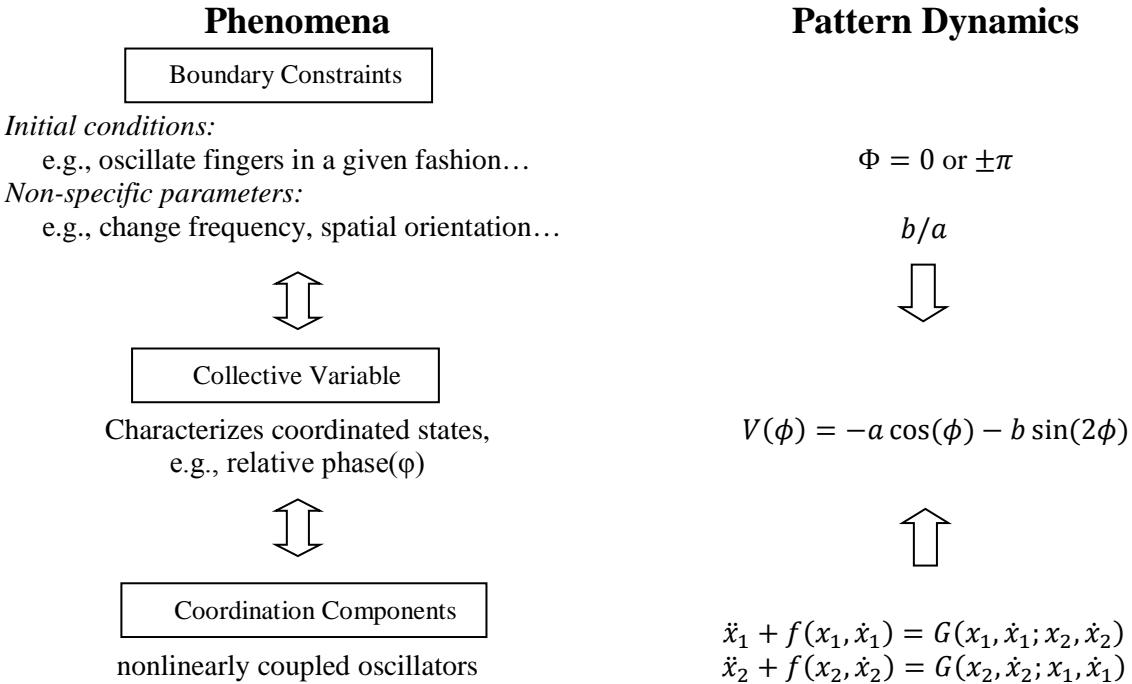


Figure 1: Tripartite Scheme for rhythmic finger coordination. Adapted from Kelso (1995).

provided by the collective variable is a low dimensional description of the behavior of a high dimensional dynamic system. While the system of interest may be highly complex, the behaviors of the system will be modeled by the collective variable in terms of relatively few patterns, determined by the energy requirements to maintain those patterns. The collective variable modeled as the relative phase of the left and right fingers:

$$V(\phi) = -a \cos(\phi) - b \sin(2\phi)$$

describes bimanual coordination as a function of the energy required ($V(\phi)$) to maintain coordination at relative phase (ϕ). This equation predicts critical slowing near phase transitions and hysteresis effects that are observed in bimanual coordination tasks. In doing so, it provides a model of the system's behavior in terms of its energy dynamics; the system settles into low energy behaviors unless it is forced to do otherwise. Significantly, while a collective variable typically characterizes relatively large scale/slower features of system, they constrain the behavior of smaller scale/faster features of the system. In synergetics, this is called an 'enslaving principle' and plays an important role in explanation of certain types of systems. It is important to note that neo-mechanists often overlook this and other features of the theory underlying the use of dynamic systems models for explanation.

The *coordination components* are defined by sets of coupled differential equations that describe the energetic constraints on the components that are coordinated. When the boundary constraints define the system as a coupled pair of oscillators, the coordination components model the energetic constraints governing the individual oscillators

and their coupling. In the HKB model, the coordination components are individually oscillating fingers, described as nonlinearly coupled oscillators. It is not possible to accurately define the behavior of either component independently of the other in the context of the boundary constraints because of the coupling between the oscillators. In traditional mechanistic terms (Wimsatt, 2007; Bechtel and Richardson, 2010), the system is at best minimally decomposable, if not non-decomposable, which is to say that it is not subject to neo-mechanistic explanation. The behavior of each component is determined as strongly, or more strongly, by the state of other coordination components that make up the system than by its own intrinsic dynamics. The use of the term ‘component’ in this case is not identical to its use in neo-mechanistic descriptions of systems. The individual equations and variables that describe the coordination components do not map onto neo-mechanistically defined lower level components of a system. Instead, the variables and parameters are defined in terms of energy expenditure. Neo-mechanists, of course, do not claim that HKB is a neo-mechanistic equation. Rather, they would say that it is a mere description of data.

Applying the Scheme to the Brain

Once several models of a system are developed at different scales of description, the models can be treated as coordination components for new scales of description (Potochnik and McGill, 2012). In the context of dynamic systems research the goal of comparing and coordinating models is to discover the extent to which a system's

Phenomena

Boundary Constraints

Initial conditions:

e.g. Neural field under periodic stimulation
during rhythmic motor activity...

Non-specific parameters:

e.g. change frequency, spatial mode, neural sheet dynamics, ...



Collective Variable

Characterizes coordinated states,
e.g., relative phase(ϕ)



Coordination Components

Nonlinearly coupled oscillators,
e.g. Neural sheet amplitude with sensory
and auditory inhomogeneities

Pattern Dynamics

$\Phi = 0$ or $\pm\pi$

$\omega_i/2\Omega$, etc...



$$V(\phi) = -a \frac{y_{st}^2}{x_{st}} \cos(\phi) - b \cos(2\phi)$$



$$\begin{aligned} x_{10} &= \left(-\frac{1}{2}(y_0 + y_1) + \frac{1}{4}\epsilon \cos 2\phi \right) x_{10} - B \left(\frac{1}{2}x_{10}^2 + \right. \\ &\quad \left. 2(\tilde{\epsilon} - \epsilon_0^2)y_{20}^2 \right) x_{-10} \\ y_{20} &= \\ &\quad \left(-\frac{1}{2}y_0 + \frac{1}{4}\epsilon \epsilon_0 (1 + \cos 2\theta) - \frac{\epsilon \epsilon_0 \tilde{\omega}_0}{8\Omega} \sin 2\theta \right) y_{20} - \\ &\quad B \left(\left(1 + \frac{\epsilon_0}{2} \right) x_{10}^2 + \frac{2}{2} (1 + 2\epsilon_0^2) y_{20}^2 \right) y_{20} \end{aligned}$$

Figure 2: Tripartite Scheme for an application of the neural field equation to rhythmic coordination experiments (Jirsa and Haken, 1997; adapted from Kelso, 1995).

behavior is governed by the energetic constraints of the system at different scales of description. Notably, this is where dynamical models gain their explanatory force (Kaplan and Craver, 2011). Related models, i.e. models that are related to the same phenomena, are examined to determine if there are boundary constraints that can be applied to the coordination components of both models to produce equivalent collective variables. It must then be shown that the selected boundary constraints can be applied within the models such that they consistently map onto the system of interest. When a system's behavior(s) can be modeled at a variety of scales such that each model has the same collective variable, we can take this as evidence that across these scales the system's behavior is governed by the same energetic constraints. This means that regardless of the small-scale physical stuff that makes up the system, an explanation of the system's behavior must include an account of the energetic constraints or enslaving principles that the system is subject to. This would be a clear case of a thoroughly dynamical system. Representing such a system in neo-mechanistic terms may be useful for some types of interventions on the system, but it wouldn't necessarily add to our understanding of why the system does what it does. In fact, since such a system is at best minimally decomposable, any neo-mechanistic explanation of the system's behavior would be lossy, an over-simplification of

the dynamic systems model already provided. This is because a mechanistic model necessarily ignores the effects of non-linear coupling both in defining the system's components and defining the system's behavioral modes.

We have, so far, focused on the structure of a dynamic systems approach to modeling and explaining systems of interest. We can now apply this understanding to recent work relating the neural field model to rhythmic coordination behaviors. A number of studies have shown that the dynamic models that describe various scales of activity related to rhythmic finger movements are equivalent (Jantzen et al, 2009; Jirsa and Haken, 1996; Jirsa and Haken, 1997; Jirsa et al, 1998). V.K. Jirsa and H. Haken developed a dynamic systems model of the electromagnetic activity in the brain called the neural field model (Jirsa and Haken, 1996; Jirsa and Haken, 1997; Jirsa et al, 1998). Recently, neo-mechanists (Kaplan and Bechtel 2011; Kaplan and Craver 2011) argue that the neural field model is mechanistic. However it should be clear that it is not neo-mechanistic from a careful reading of the model in the context of the Tripartite Scheme. Further, Jirsa and Haken make this clear when they state, "The [collective variables] are determined and created by the cooperation of microscopic quantities, but at the same time the [collective variables] govern the behavior of the whole system" (Jirsa and Haken, 1996, p960). That is, the behavior of the

components alone is not responsible for and does not explain the behavior of the system. Instead, the components make up the system, and the system constrains the behavior of the components. The neural field model is a dynamic systems model of pulse to wave and wave to pulse behaviors of neural populations (See Figure 2). In the study we will look at, one cited by Kaplan and Bechtel and Kaplan and Craver, the neural field model is applied to the neocortex, which is modeled as a medium represented by a one dimensional closed surface (Jirsa et al, 1997). This medium (neural sheet) consists of neural ensembles characterized by two state variables, waves of extracellular fields and intracellular fields that correspond to pulses. The field model is thus a model of energy transfer and conversion through a medium. The model's boundary constraints are defined in terms of excitatory and inhibitory pulse and wave amplitudes (Jirsa and Haken, 1996). The parameterization of the model is based on the electrical properties of the neuron ensembles as well as their spatial scale and connectedness in the neocortex. The coordinated components are the neural sheet and inhomogeneities in the neural sheet. The inhomogeneities are the result of neural ensembles coupled to brain and body areas outside the neural sheet and map on to traditional functional brain areas. The coupling between the neural sheet and inhomogeneities is accomplished through the neural field equation which is derived from the wave-pulse and pulse-wave conversions that take place within neural sheet. When the coordination components are coupled by the neural field equation, the system's mode can be calculated. The collective variable is defined in terms of the energy distribution patterns of the system (Jirsa et al, 1998; Jirsa and Haken, 1997).³ The energy distribution patterns can be thought of as the pattern of ripples in a puddle at a particular moment. According to the neural field model, the behavior of the system is explained and described by stable distribution patterns of wave and pulse fields in the neural sheet. Thus, the neural field model is a model of the energy dynamics of the system. At the neural population scale of description, the neural field model explains the system's behavior according to the distribution of energy throughout the system.

In studies by Jirsa, Fuchs, and Kelso (1998) and Jirsa and Haken (1996) the neural field model is applied to neural imaging work done in rhythmic finger coordination experiments. A participant is subjected to an auditory stimulus while asked to perform a rhythmic tapping task. During the study, the participant's neural activity is recorded using a SQUID array. In these studies three inhomogeneities are modeled in the neural sheet. These map onto the auditory cortex, and the motor and sensorimotor brain areas. The sensorimotor area is defined as a subset neural population in the neural sheet that is driven, in part, by the movements of the finger it is coupled to. The motor area is defined as a subset neural population in the neural sheet whose behavior drives finger movements. The

auditory cortex is defined as a subset neural population that is driven by an acoustic stimulus. Despite the reference to traditional functional brain areas, it is important to recognize that these areas do not play a neo-mechanistic role in the system. Due to the structure of the neural field model, none of the inhomogeneities can be adequately defined independent of the neural sheet it is imbedded in, including the states of the other inhomogeneities. Thus, they are at best minimally decomposable within the model, which is insufficient for neo-mechanistic explanation. An analogy may help. Returning to the puddle, if we attempt to model three drops of water falling simultaneously into a small puddle of water, the ripple patterns produced by any individual drop could not be modeled without also modeling the other drop and the relevant properties of the puddle. The neural field model represents the neural sheet and functional brain areas in a similar way. Significantly, within certain boundary constraints, the system's behavior isn't characterized by the behavior of any particular inhomogeneity or aspect of the neural sheet. Its behavior is characterized by the energy distribution patterns of the neural sheet with embedded inhomogeneities. Stretching the metaphor, the behavior isn't characterized by the individual drops of water; rather it's characterized by the pattern of waves in the puddle.

Put back in terms of cognitive science, there is no sense in which the activity of the neural field produces, explains, determines or accounts for the rhythmic tapping behavior. Instead, the boundary conditions determine the tapping behavior, which determines the coordination pattern described by the neural field model. The neural field theory is an application of the Tripartite Scheme, in which the tapping task is a boundary condition, the collective variable models cortical coordination patterns, and the coordination components are coupled inhomogeneities. The neural field theory is a dynamical model just as the HKB model is. The fact that it is, in part, a dynamical model of the brain does not make it a neo-mechanistic model. As noted above, the application of a model to a physical system is insufficient to make it a neo-mechanistic model. The model must also be decomposable according to the neo-mechanist account. If a neo-mechanist wishes to discard the condition of decomposability, then she does so at the cost of discarding the feature of neo-mechanistic explanations that makes them distinct from more general accounts of naturalistic explanation. In a general account of this type a researcher develops a model of a system's behavior and maps that model onto a physical system without remainder. We take it that our discussion of the Tripartite Scheme and the neural field model has gone at least some distance towards showing that dynamic systems explanations fit this more general account of explanation. The explanation of the phenomena of interest is embodied in the model and not the application of the model to some set of components and relations in a system of interest. However, we acknowledge that much more work needs to be done, and in part has been done, to show that this is sufficient to justify the claim that

³ This is referred to in the research as the system's "spatial mode."

dynamic system's explanations really are explanatory without reference to neo-mechanist methods.

Applying the neural field model to brains does not require adopting a neo-mechanistic account of explanation. In fact, once the neural field model is applied at the neural ensemble scales and the Tripartite Scheme is elaborated, the results can be compared and coordinated with other scales of description. This is what Kelso and his colleagues have been up to for the past couple of decades. The example just given has been a part of a larger body of work to show that the collective variable for rhythmic coordination behavior at the neural ensemble scale is equivalent to the collective variable for rhythmic coordination behavior at the finger behavior scale (Jantzen et al, 2009; Jirsa and Haken, 1996; Jirsa and Haken, 1997; Jirsa et al, 1998). Doing this is a part of determining the extent to which the rhythmic coordination behaviors depend on energy dynamics. This is distinct from the neo-mechanist project of determining the extent to which rhythmic coordination behaviors depend on discrete input and output control of hierarchically ordered and discrete components. Given the success of Kelso and his colleagues with respect to this project, the dynamics systems position is increasingly plausible (Jantzen et al, 2009; Jirsa and Haken, 1996; Jirsa and Haken, 1997; Jirsa et al, 1998). More importantly for current purposes, the SFK argument loses its force entirely. It is simply inaccurate to claim, as Kaplan and Bechtel and Kaplan and Craver do, that recent work on the neural field by Kelso and colleagues is mechanistic in character. Scott Freaking Kelso has been publishing articles about the brain for nearly two decades at this point, but explaining features of the brain is not identical to explaining neo-mechanistically.

Conclusion

In this paper, we have sought to set straight some misconceptions about the neural field theory. In particular, we have sought to clarify the role of the Tripartite Scheme in structuring the methodology and investigative questions of dynamic systems researchers. Contrary to the claims of some neo-mechanists, dynamic systems researchers have not had resort to neo-mechanistic methodologies or explanatory strategies. Instead, dynamic systems researchers have continued to engage and study cognitive systems using the methodologies and explanatory strategies of dynamic systems theory. This has resulted in research that spans and connects scales of descriptions and that defies traditional reductive and neo-mechanistic understandings. Kelso and his colleagues have continued to expand their understanding of the role of dynamic systems theory in cognitive systems over the years and they have done so without relying on neo-mechanistic models.

References

- Bechtel, W. (2009). Constructing a philosophy of science of cognitive science. *Topics in Cognitive Science*, 1(3), 548–569.
- Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 36(2), 421–441.
- Bechtel, W., & Abrahamsen, A. (2010). Dynamic mechanistic explanation: Computational modeling of circadian rhythms as an exemplar for cognitive science. *Studies in History and Philosophy of Science Part A*, 41(3), 321–333.
- Bechtel, W., & Richardson, R. C. (2010). *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research* (Reissue.). The MIT Press.
- Haken, H., Kelso, J. A. S., & Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological cybernetics*, 51(5), 347–356.
- Jantzen, K. J., Steinberg, F. L., & Kelso, J. A. S. (2009). Coordination dynamics of large-scale neural circuitry underlying rhythmic sensorimotor behavior. *Journal of cognitive neuroscience*, 21(12), 2420–2433.
- Jirsa, V. K., & Haken, H. (1996). Field theory of electromagnetic brain activity. *Physical Review Letters*, 77(5), 960–963.
- Jirsa, V. K., & Haken, H. (1997). A derivation of a macroscopic field theory of the brain from the quasi-microscopic neural dynamics. *Physica D: Nonlinear Phenomena*, 99(4), 503–526.
- Jirsa, V. K., Fuchs, A., & Kelso, J. A. S. (1998). Connecting cortical and behavioral dynamics: bimanual coordination. *Neural Computation*, 10(8), 2019–2045.
- Kaplan, D. M., & Bechtel, W. (2011). Dynamical Models: An Alternative or Complement to Mechanistic Explanations? *Topics in Cognitive Science*, 3(2), 438–444.
- Kaplan, D. M., & Craver, C. F. (2011). The Explanatory Force of Dynamical and Mathematical Models in Neuroscience: A Mechanistic Perspective. *Philosophy of science*, 78(4), 601–627.
- Kelso, J. A. S. (1995). *Dynamic patterns: the self-organization of brain and behavior*. Cambridge, Mass: MIT Press.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of science*, 1–25.
- Potochnik, A., & McGill, B. (2012). The Limitations of Hierarchical Organization. *Philosophy of Science*, 79(1), 120–140.
- Stepp, N., Chemero, A., & Turvey, M. T. (2011). Philosophy for the rest of cognitive science. *Topics in Cognitive Science*, 3(2), 425–437.
- Wimsatt, W. C. (2007). *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Harvard University Press.