Mechanism and Explanation in Cognitive Neuroscience

Barbara Von Eckardt and Jeffrey S. Poland†‡

The aim of this paper is to examine the usefulness of the Machamer, Darden, and Craver (2000) mechanism approach to gaining an understanding of explanation in cognitive neuroscience. We argue that although the mechanism approach can capture many aspects of explanation in cognitive neuroscience, it cannot capture everything. In particular, it cannot completely capture all aspects of the content and significance of mental representations or the evaluative features constitutive of psychopathology.

1. Introduction. This paper has two starting points: recent work in philosophy of science on the notion of mechanism and the newly emergent field of cognitive neuroscience. We have several aims in bringing together this particular area of philosophy of science and this particular piece of science. Our primary aim is to examine the usefulness of the mechanism approach to gain an understanding of explanation in cognitive neuroscience. Although most research on the mechanism approach has involved using it as a tool to reconstruct aspects of the biological sciences at a fairly low level, such as molecular biology and neurobiology, there is work, specifically Bechtel and Richardson 1993 and Craver and Darden 2001, that extends the approach to the higher reaches of neuroscience, and, in particular, to parts of neuroscience that border on psychology. It is thus quite appropriate to ask whether the mechanism approach can be useful in elucidating cognitive neuroscience. In addition, we have two secondary aims: to draw some connections between the mechanism approach and work in philosophy of cognitive science and to expand the discussion of explanation in the mind/brain sciences to include consideration of psychopathology.

†To contact either author, please write to: Division of Liberal Arts, 2 College Street, Rhode Island School of Design, Providence, RI 02903-2784; e-mail: bvonecka@risd.edu or jpoland@risd.edu.

‡We would like thank Carl Craver for organizing the symposium session at which this paper was originally presented and Carl and Lindley Darden for useful discussion.

Philosophy of Science, 71 (December 2004) pp. 972–984. 0031-8248/2004/7105-0029$10.00

We will be arguing for two claims:

1. The mechanism approach can capture many aspects of cognitive neuroscience.
2. However, there are aspects of cognitive neuroscience that it cannot capture, specifically, certain features of representations and certain aspects of psychopathology.

2. The Mechanism Approach and Cognitive Neuroscience.

2.1. The Mechanism Approach. In the past decade, a number of philosophers of science have begun to develop an approach to the nature of models, explanation, and discovery based on the notion of a mechanism (Bechtel and Richardson 1993; Glennan 1996; Machamer, Darden, and Craver 2000; Craver 2001; Craver and Darden 2001). A mechanism is defined as "entities and activities organized such that they are productive of regular change from start or set-up to finish or termination conditions" (Machamer, Darden, and Craver 2000, 3). The setup conditions are entities and their properties (e.g., parts, structure, spatial relations, and orientations) and enabling conditions (e.g., available energy, pH, electrical-charge distributions). Termination conditions are idealized states or parameters that constitute a privileged endpoint (e.g., rest, equilibrium, elimination of something, production of something). The intermediate activities that take the mechanism from the setup conditions to the termination conditions are intervening entities and activities (e.g., causal chains, forks, joins, cycles) (Machamer, Darden, and Craver 2000).

Mechanisms are characterized by both temporal and spatial organization. With respect to temporal organization, a mechanism, typically, has stages that are ordered and that have characteristic rates and durations. With respect to spatial organization, the stages of a mechanism are typically localized with respect to particular entities and their activities, and are connected to each other in various ways. These connections often depend on the structure of the relevant entities and their spatial interrelations. Mechanisms can be described with varying degrees of completeness. A complete mechanistic description will exhibit "productive continuity" in terms of "bottoming out" lower-level mechanisms, where bottoming-out mechanisms are "the components that are accepted as relatively fundamental or taken to be unproblematic for the purposes of a given scientist, research group, or field" (Machamer, Darden, and Craver 2000, 13). In contrast, a *mechanism schema* will be "a truncated abstract description of a mechanism that can be filled with descriptions of known component parts and activities" (2000, 13), whereas a *mechanism sketch*

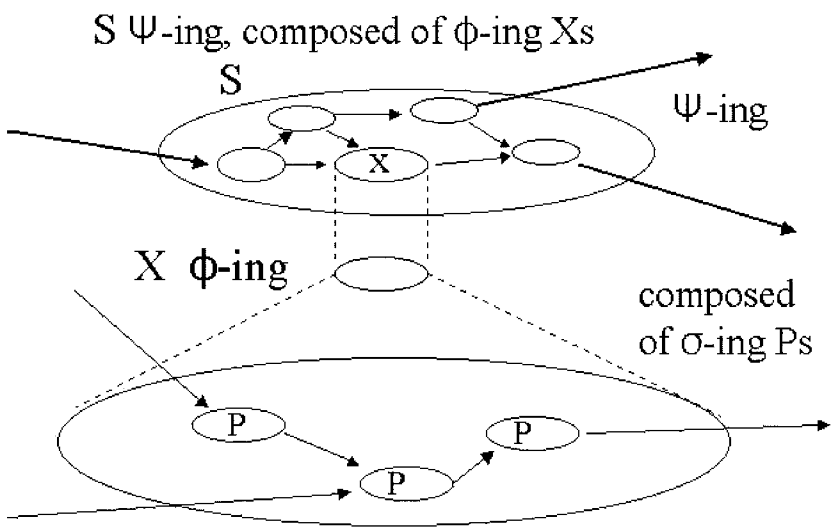

Figure 1. Mechanism hierarchy (adapted from Craver 2001).

is "an abstraction for which bottom out entities and activities cannot (yet) be supplied or which contain gaps in the stages" (2000, 18).

Mechanism descriptions vary not only with respect to degree of completeness or specificity, they can also vary with respect to their "vertical" scope. To capture the idea that a mechanism can consist of other mechanisms, Craver has developed the notion of a *mechanism hierarchy*. In Figure 1, we see a mechanism *S*, which is Ψ-ing, composed of smaller entity *X*s, which are Φ-ing. Each of these *X*s is itself a mechanism consisting of even smaller entity *P*s, which are σ-ing (Craver 2001).

Mechanistic descriptions can be used for various explanatory purposes. Craver distinguishes three kinds of mechanistic explanations: etiological, constitutive, and contextual. Only the second—constitutive—is relevant for our purposes. The explanandum in a constitutive explanation is the activity (or "phenomenon") exhibited by a mechanism, taken as a whole or, more precisely (if one adopts the view that explananda are always *questions*, which we do), how (actually, possibly, or plausibly) *S* exhibits the activity (process, capacity) of Φ-ing. The explanans is then, simply, a description of the internal mechanistic structure of *S*, in terms of its lower-level entities and activities and their interrelations, in virtue of which *S* does, or can, or possibly could, or plausibly could Φ. In other words, a constitutive mechanistic explanation of how *S* Φs makes reference to the mechanism that constitutes *S* and its Φ-ing (Craver 2001).

A similar explanatory notion has become standard in the philosophy of psychology where the "phenomenon" in question is usually taken to be a capacity rather than an activity. Cummins suggests that a system's or organism's capacity to do something can be explained by a *functional analysis* of that capacity. By a functional analysis of some capacity *C*, he means, roughly, a breakdown of the complex capacity into a set of constituent capacities ($c_1, c_2, \ldots, c_n$), with a specification of the sequence in which those constituent capacities must be exercised for the result to be a manifestation of the complex capacity (Cummins 1975). Von Eckardt points out that a functional analysis for a complex capacity *C* only provides us with a *possible* explanation of how the organism has *C*. The reason is that most complex capacities can be broken down into various *alternative* sequences of constituent capacities, the exercise of each of which is sufficient for the exercise of the complex capacity. What is required for an actual explanation is not only a functional analysis but some reason to believe that the system or organism in question actually has (or exercises) *C* in virtue of actually having (and exercising) the constituent capacities of the functional analysis in the order specified by the functional analysis, namely, the functional analysis must be *structurally adequate* (1977). Von Eckardt further suggests the following as sufficient conditions for structural adequacy:

> A functional analysis A of how an organism O has capacity C is structurally adequate if there exists a *structural decomposition* of O into component (physical) parts such that
>
> (1) Operation of those parts results in O's exercising C.
> (2) For each constituent capacity of A, there exists at least one of those structural components that has that capacity.
> (3) The order in which those component parts operate when O is exercising C mirrors the algorithm specified in A.
>
> Condition (1) guarantees the existence of a physical mechanism underlying the capacity C. Condition (2) ensures that some part of the mechanism corresponds to each constituent capacity specified by the functional analysis. Condition (3) guarantees further that the sequence of operation of these physical parts corresponds to that specified by the functional analysis. (1977, 41)

The model that describes a possible structural decomposition of *O* such that if it were true of *O*, conditions (1)–(3) would be met, is a *functional component model of how* O *has* C.

Note that all of this is easily translated into the language of the mechanism approach. (Indeed, the breakdown of *O* into component [physical] parts whose operation results in *O*'s exercising *C* is explicitly referred to in Von Eckardt (1977) as a "mechanism.") A functional analysis can be viewed as a mechanism sketch that begins to provide a constitutive mechanistic explanation of how *O* exercises *C*. It is only a mechanism sketch because constituent capacities are hypothesized without reference to the constituent entities that (presumably) have those capacities. However, a functional analysis combined with a functional component model is equivalent to a complete constitutive mechanistic explanation. Both can be further specified with respect to modality or epistemic strength, and for both we can distinguish an explanation of a capacity itself vs. an explanation of the *exercise* of a capacity.

2.2. Cognitive Neuroscience. In the 1970s and 1980s, there were two multidisciplinary fields studying the mind/brain without much interaction: cognitive science and the neurosciences. As Bechtel has emphasized, both acknowledged the importance of the other in principle, but for both there were reasons for not taking research in the other all that seriously. On the one hand, cognitive science paid lip service to the importance of neural realization but there was no clear picture, even in principle, of how classical ("symbolic") computational models could be realized in the brain and the data coming out of neuropsychology and neurology was not very useful in constraining those models. On the other hand, the major successes in the neurosciences had been at the cellular and molecular levels but there were no techniques for studying higher cognitive processes other than the classical neurological and neuropsychological ones (Bechtel 2001).

All that changed in the 1980s with the emergence of the technologies of PET, MRI, and fMRI. Within a decade there were eight major centers for cognitive neuroscience funded by the McDonnell Foundation and the Pew Charitable Trust, a new *Journal of Cognitive Neuroscience*, and a new Cognitive Neuroscience Society. The aims of this new field were, first, to further the development of adequate psychological models of cognition by generating and bringing to bear a variety of kinds of neural data and, second, to develop adequate hypotheses about the neural localization and realization of cognitive structures and processes described psychologically. Michael Gazzaniga, the first editor of the *Journal of Cognitive Neuroscience*, wrote in his opening editorial that the journal was for cognitive scientists who "now believe that it is maximally fruitful to propose models of cognitive processes that can be assessed in neurobiologic terms" and neuroscientists who believe they must "actually come to grips with the complexities of psychological processes involved in any particular mental capacity" (Gazzaniga 1989). Something very similar could be said of the field itself.

Another way to understand the explanatory aims of cognitive neuroscience is provided in Posner and Rachle's nice introduction to the field, *Image and Mind* (1994). There they distinguish five levels for linking the cognitive and neural levels of analysis—the cognitive system, the mental operation, the performance domain (or pathway activation), the neural system, and the cell—each of which is associated with its own methods of investigation. For our purposes, what is important to note is that the typical cognitive neural model includes an information processing story of the capacity under study (usually, at a fairly crude level) plus hypotheses about the functional localization of the constituent capacities and operations posited by that story. The hope is, then, that in the future, the model will be supplemented by hypotheses detailing the neural implementation of these operations.

3. Cognitive Neuroscience through the Lens of the Mechanism Approach. Does the mechanism approach have the resources to adequately "capture" the field of cognitive neuroscience? We shall explore this question by looking at Craver and Darden's "Discovering Mechanisms in Neurobiology: The Case of Spatial Memory" (2001), the mechanism approach paper that comes closest to a reconstruction of cognitive neuroscience.

The focus of Craver and Darden's paper is recent investigations of spatial memory in the mouse. What interests them is that this research provides a wonderful example of a multilevel hierarchy of mechanisms. The phenomenon under investigation is the ability of the mouse "to learn to navigate through a novel environment." It is this ability that is labeled "spatial memory." Underlying this ability is said to be the generation of a spatial map that is localized in the hippocampus of the mouse brain. This hippocampal capacity is, in turn, claimed to involve long-term potentiation (LTP) of synapses in the hippocampus which, in turn, is currently understood as involving NMDA receptor activity.

Note that three of the levels described in the Craver and Darden analysis of spatial memory research correspond nicely with the levels of Posner and Rachle's hierarchy and with the levels typically distinguished in philosophy of cognitive science (see Table 1). "Spatial memory" is, in fact, a label for a complex capacity or what Posner and Raichle call a "cognitive system." Underlying this capacity or system is the mental operation or information-processing (constituent) capacity of generating a spatial map. This is, then, localized in a neural system and given a partial explanation in terms of the cellular-level activity of long-term potentiation.

The most important aspect of Craver and Darden's discussion of this case concerns the relationship of these various levels to one another. They write:

> The levels in this sort of hierarchy stand in part-whole relations to

TABLE 1. Comparison of the Posner-Raichle and Philosophy of Cognitive Science Frameworks.

Level	Example	Philosophy of Cognitive Science
Cognitive system	Language, attention, motor control	Complex cognitive capacities
Mental operation	Next, rotate, zoom	Component capacities
Performance domain (pathway activation)	Facilitate, inhibit	Information processing model (computational processes over representations)
Neural system	Parietal lobes	Functional localization of component capacities
Cell	Primary visual area	Neural implementation of information processing model

> one another with the important additional restriction that the lower-level entities and activities are components of the higher-level mechanism. The binding of glutamate to the NMDA receptor is a lower-level activity in the mechanism of LTP, and LTP is thought to be a lower-level activity in spatial map formation, which, of course, is thought to be an activity in the mechanism of spatial memory. (Craver and Darden 2001, 117)

This quote plus our previous discussion of the mechanism approach suggests a number of claims about the commitments of the approach:

> M1. All mechanism hierarchies are part-whole hierarchies.
>
> M2. A system constitutes a mechanism only if it is part of a (part-whole) mechanism hierarchy.
>
> M3. Constitutive mechanistic explanations make reference to mechanisms.
>
> M4. An explanation counts as a constitutive mechanistic explanation of phenomenon *P* only if it is possible, in principle, to provide a multilevel part-whole constitutive mechanistic explanation of *P*.

Although advocates of the mechanism approach have not directly addressed questions concerning explanation in cognitive neuroscience, Craver and Darden's paper strongly *suggests* the attitude they might take: that cognitive neuroscience constitutive explanations of human cognitive capacities *are nothing but constitutive mechanistic explanations*. This extrapolation of Craver and Darden's treatment of spatial memory research is reasonable because it is clear, on their view, that neurobiology is attempting to provide a constitutive mechanistic explanation of spatial memory and spatial memory is precisely the kind of phenomenon cognitive neuroscience is attempting to explain (albeit in humans instead of mice).

4. Why Constitutive Explanation in Cognitive Neuroscience Is More than Mechanistic Explanation.

4.1. Constitutive Explanation in Cognitive Neuroscience Is Representational. As we indicated earlier, a type of explanation, namely, functional analysis, similar to constitutive mechanistic explanation, has been part of the philosophical reconstruction of cognitive science for more than twenty-five years. It is similar because it seeks to explain certain global properties of a system or organism in terms of that system's or organism's constitution. More specifically, it seeks to explain how organism *O* has such-and-such properties (including capacities) by reference to properties (including subcapacities) of parts of *O*. It has been understood for many years, however, that constitutive explanation in cognitive science is a particular *species* of constitutive explanation, namely, it is a type of constitutive explanation that involves appeal to *representations*. In particular, the basic *explananda* of cognitive science are the human cognitive capacities along with a host of properties of those capacities that involve intentional states (e.g., pragmatic evaluability, coherence, and reliability). To explain the intentionality of mental states, cognitive scientists have adopted the strategy of "semantic heritability." Roughly speaking, the strategy is to explain intentionality by reference to representations with content and significance for their "owner." Given this strategy, the explanation of how a person exercises a certain capacity (e.g., reads a sentence) appeals to a computational story (either "symbolic" or connectionist) that details the steps or stages the person's mind/brain goes through as it goes from a visual (or auditory) representation of an inscription (or utterance) of the sentence to a representation of its meaning, that is, the representational state underlying our "understanding" of the sentence. The explanation of how a person *has* (rather than exercises) a certain capacity appeals to the computational and representational resources required for any exercise of that capacity. Both kinds of explanation are constitutive and both involve representations.

What does all of this have to do with cognitive neuroscience? If theoretical models in cognitive neuroscience are information-processing models mapped onto the anatomy and physiology of the brain, then if constitutive explanations in cognitive science involve representations, then so will constitutive explanations in cognitive neuroscience. And, indeed, this is precisely what we find, although sometimes the hypotheses are framed in terms of intentional capacities rather than strictly in terms of representations. For example, it has been known for some time that in monkeys there are two main streams of projections from primary visual cortex to other cortical regions (Ungerleider and Mishkin 1982): a *ventral* stream from V1 to the infero-temporal cortex and a *dorsal* stream from V1 to

https://doi.org/10.1086/425946 Published online by Cambridge University Press

posterior parietal cortex. There has been some dispute, however, about the function of these two streams. Ungerleider and Mishkin's hypothesis was that the ventral stream played a major role in the visual identification of objects, whereas the dorsal stream concerned the localizing of objects in visual space. However, Goodale and Milner have argued that the important functional distinction is between visual perception and the visual control of action (Goodale and Milner 1992; Milner and Goodale 1993). According to this account,

> both streams process information about object characteristics, such as size, orientation, and shape, and both process information about spatial location. Each stream, however, uses this visual information in different ways. Transformations carried out in the ventral stream permit the formation of perceptual and cognitive representations that embody the enduring characteristics of objects and their spatial relations with each other; those carried out in the dorsal stream, which utilize instantaneous object features that are organized within egocentric frames of reference, mediate the control of goal-directed actions. (Goodale, Jakobson, and Servos 1996, 107)

4.2. Mental Representations. We have suggested that cognitive neuroscience, like cognitive science itself, invokes mental representations to explain cognition. What is it about resorting to representations that poses problems for the mechanism approach? To answer this question, we must consider what mental representations are, with a focus on how they are constituted.

Mental representations are best understood in terms of Charles Peirce's general triadic theory of representation. They consist of: (1) a representation bearer, which has some sort of (2) content, which, in turn, has (3) significance for an interpreter who, in this case, is the person in whose mind/brain the representation resides. Insofar as a mental representation has these three aspects, a constitutive account of mental representation must explain *in virtue of what* any given mental representation has (1), (2), and (3). A constitutive explanation of a mental representation's bearer is the most straightforward. If one views the mind/brain as not only a representational system but also a computational one, then the bearers of mental representations will be nonsemantic, purely formal, computational structures or states. Precisely what kinds of computational structures or states will depend on whether one views the mind-brain as a "classical" computational system, as a "connectionist" system, or as some other kind of computer. For example, if the mind/brain is viewed as a classical computational system, then the representation bearers of mental representations will be data structures. If it is viewed as a connectionist computer, the representation bearers will be patterns of activation (for recurrent representational states) or patterns of weights (for stored representations). All of this is, obviously, quite compatible with a mechanism approach.

In contrast, difficulties for the mechanism approach emerge when we consider how both the content and significance of mental representations might be realized. Consider first, content. Suppose we are trying to explain how people exercise the capacity to provide a use verb (e.g., 'hit') in response to a visually presented name of an object (e.g., 'hammer'). A cognitive psychologist or cognitive neuroscientist will most likely hypothesize that the early stage of exercising this capacity is to form a representation of the stimulus word. Let's suppose further that the representation will be an orthographic one, that is, it will represent the word 'hammer' in a way that is similar to a standard printed representation, by a sequence of component representations of the letters that spell the word 'hammer'.

A major foundational question of cognitive science (and, *ipso facto*, cognitive neuroscience, insofar as it incorporates the information-processing level of description) is how content is realized naturalistically. Philosophers (and a few psychologists) have proposed a variety of theories of content determination (see Von Eckardt 1993 for an overview) the details of which need not concern us here. The important point is that many of these theories are *externalistic*. That is, the naturalistic properties and relations of a representation bearer that realize its content are taken to be properties and relations that extend beyond the head, and, indeed, in the case of the evolutionary version of the teleosemantic theory of content determination, into the distant past. Furthermore, there is every reason to believe that the correct theory of content determination will have, at a minimum, an externalistic component. Clearly, such externalistic components will not be part of any neural mechanism.

What about significance? The significance of a mental representation is how it is used by its "owner." Von Eckardt (1993) argues that within cognitive science (and so, again *ipso facto*, within cognitive neuroscience) the significance of a representation for some subject S consists in the set of all possible determinate computational processes contingent upon entertaining (activating) that representation in *S*. In other words, it consists of something like downstream "narrow" functional role. Since the computational consequences of entertaining a representation are "in the head," it might not be immediately obvious why this account of representational significance would cause problems for the mechanism approach to constitutive explanation. The reason turns on the fact that any given (fairly specific) type of representation, say, the orthographic representation of the word 'hammer' mentioned above, will play a role in many capacities and, hence, will have a wide variety of computational consequences. However, when we are delineating the mechanism underlying a specific capacity, say, reading, we will be focusing on only *some* of those consequences. So there are likely to be other consequences that, on the one hand, constitute part of the *significance* of the representations that are part of the mechanism of reading, but that, on the other, play no role in the mechanism itself.

A similar challenge to a purely mechanistic reconstruction of cognitive neuroscience arises in the domain of psychopathology. For example, according to the *Diagnostic and Statistical Manual of Mental Disorders*, 4th edition (American Psychiatric Association 1994), a delusion is "a false belief based on incorrect inference about external reality that is firmly sustained despite what almost everyone else believes and despite what constitutes incontrovertible and obvious proof or evidence to the contrary" (1994, 765). So in order to explain, at the psychological level, in virtue of what people have delusions, cognitive neuroscientists will have to resort to the content of beliefs and, hence, to representations. But, as we have seen, a mechanism approach that embraces M1–M4 will only be able to capture the formal aspects of the relevant representations, not all of their content nor all of their significance.

Psychopathology introduces another feature that raises problems for a mechanistic reconstruction of cognitive neuroscience. This feature is already introduced by our example of delusion, namely, the fact that many psychopathological phenomena are classified *normatively.* A variety of norms are referenced in the classification of psychopathology, some based on statistics or natural design ("abnormal"), some that are social or prescriptive ("bizarre," "inappropriate," "deviant") and some that are epistemic ("unreasonable fear," "poor insight," "nonresponsiveness to evidence"). The point is that if we are interested in providing a constitutive explanation of some piece of psychopathology, part of the explanation must include an account as to why the cognition or behavior merits the evaluative label. But such an account will involve a comparison of the cognition or behavior with a norm and determine whether the norm is statistical, based on natural design, based on community standards, or epistemic. Any such account will extend well beyond what one would naturally assume to be the mechanism of (or the breakdown of the mechanism of) the cognition or behavior in question.

In sum, then, there are several reasons why constitutive explanation in cognitive neuroscience is more than mechanistic explanation and, hence, why the mechanism approach is unable to reconstruct all aspects of cognitive neuroscience. First, the constitutive explanations of human cognition provided by cognitive neuroscience make reference to representations and computational processes involving representations. But it is not possible to supply the required multilevel, part-whole constitutive mechanistic explanation of *all* aspects of any given cognitive capacity because representations are not fully constituted (realized) by lower-level neural entities or activities in the relevant part-whole mechanism hierarchy. In particular, while the neural realization of the representation bearers involved in any given capacity will be located within the relevant part-whole mechanism hierarchy, the realizations of neither the content nor the significance of those representations will be entirely so located.

Another reason that the mechanism approach cannot handle all aspects of cognitive neuroscience is that any constitutive explanation of a pathological cognition or behavior must include an explanation of the criteria by which the cognition or behavior merits an evaluative label. But such explanations require a comparison with some norm and neither the norms nor the relation to those norms will fall within the mechanism hierarchy underlying the cognition or behavior in question.

REFERENCES

American Psychiatric Association (1994), *Diagnostic and Statistical Manual of Mental Disorders*, 4th ed. Washington, DC: American Psychiatric Association Press.

Bechtel, William (2001), "Cognitive Neuroscience: Relating Neural Mechanisms and Cognition", in Peter K. Machamer, Rick Grush, and Peter McLaughlin (eds.), *Theory and Method in the Neurosciences*. Pittsburgh: University of Pittsburgh Press.

Bechtel, William, and Robert C. Richardson (1993), *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton, NJ: Princeton University Press.

Craver, Carl F. (2001), "Role Functions, Mechanisms, and Hierarchy", *Philosophy of Science* 68: 53–74.

Craver, Carl F., and Lindley Darden (2001), "Discovering Mechanisms in Neurobiology: The Case of Spatial Memory", in Peter K. Machamer, Rick Grush, and Peter McLaughlin (eds.), *Theory and Method in the Neurosciences*. Pittsburgh: University of Pittsburgh Press.

Cummins, Robert (1975), "Functional Analysis", *Journal of Philosophy* 72: 741–765.

Gazzaniga, Michael S. (1989), "Editorial", *Journal of Cognitive Neuroscience* 1: 1.

Glennan, Stuart S. (1996), "Mechanisms and the Nature of Causation", *Erkenntnis* 44: 49–71.

Goodale, Melvyn A., Lorna S. Jakobson, and Philip Servos ([1996] 2000), "The Visual Pathways Mediating Perception and Prehension", in Alan M. Wing et al. (eds.), *Hand and Brain: The Neurophysiology and Psychology of Hand Movements*. San Diego: Academic Press. Reprinted in Michael S. Gazzaniga (ed.), *Cognitive Neuroscience: A Reader*. Malden, MA: Blackwell Publishers, 106–123.

Goodale, Melvyn A., and A. David Milner (1992), "Separate Visual Pathways for Perception and Action", *Trends in Neurosciences* 15: 20–25.

Machamer, Peter, Lindley Darden, and Carl F. Craver (2000), "Thinking about Mechanisms", *Philosophy of Science* 67: 1–24.

Milner, A. David, and Melvyn A. Goodale (1993), "Visual Pathways to Perception and Action", *Progress in Brain Research* 95: 317–338.

Posner, Michael I., and Marcus E. Raichle (1994), *Images of Mind*. New York: Scientific American Library.

Ungerleider, Leslie G., and Mortimer Mishkin (1982), "Two Cortical Visual Systems", in

David J. Ingle, Melvin A. Goodale, and Richard J. W. Mansfield (eds.), *Analysis of Visual Behavior*. Cambridge, MA: MIT Press, 549–586.

Von Eckardt (Klein), Barbara (1977), "Inferring Functional Localization from Neurological Evidence", in Edward Walker (ed.), *Explorations in the Biology of Language*. Montgomery, VT: Bradford Books.

——— (1993), *What Is Cognitive Science?* Cambridge, MA: MIT Press.