How Can Philosophy Be a True Cognitive Science Discipline?

William Bechtel

Department of Philosophy, and Interdisciplinary Programs in Cognitive Science and Science Studies, University of California, San Diego

Abstract

Although philosophy has been only a minor contributor to cognitive science to date, this paper describes two projects in naturalistic philosophy of mind and one in naturalistic philosophy of science that have been pursued during the past 30 years and that can make theoretical and methodological contributions to cognitive science. First, stances on the mind–body problem (identity theory, functionalism, and heuristic identity theory) are relevant to cognitive science as it negotiates its relation to neuroscience and cognitive neuroscience. Second, analyses of mental representations address both their vehicles and their contents; new approaches to characterizing how representations have content are particularly relevant to understanding the relation of cognitive agents to their environments. Third, the recently formulated accounts of mechanistic explanation in philosophy of science both provide perspective on the explanatory project of cognitive science and may offer normative guidance to cognitive science (e.g., by providing perspective on how multiple disciplinary perspectives can be integrated in understanding a given mechanism).

Keywords: Philosophy; Mind–body problem; Mental representation; Explanation; Integrating disciplines

1. Introduction

Since the birth of cognitive science, philosophy has been included in the catalog of cognitive science disciplines. Yet it has not been a central contributor to cognitive science and its role often is not well understood. Unlike psychology (or anthropology or sociology), philosophy has neither discovered cognitive phenomena nor provided empirical evidence

Correspondence should be sent to William Bechtel, Department of Philosophy, University of California, San Diego, La Jolla, CA 92093-0119. E-mail: bill@mechanism.ucsd.edu

regarding them. Unlike artificial intelligence (AI), it has not offered computational models of cognitive activities. And unlike linguistics, it has not created theoretical frameworks, such as generative grammars, for investigating and explaining specific cognitive phenomena. Moreover, although philosophers often are engaged in discussions about mental phenomena, they seldom are members of collaborative research teams in cognitive science. They tend to keep to themselves, writing alone at their desks (or armchairs). They are less likely to attend meetings of the Cognitive Science Society and prefer philosophically oriented societies in which they are typically in the majority (e.g., the Society for Philosophy and Psychology and the European Society for Philosophy and Psychology). Likewise, they tend to publish less in *Cognitive Science* than in more recognizably philosophical journals (*Mind and Language*, *Minds and Machines*, and *Philosophical Psychology*). These considerations raise the question of whether and how philosophy can be a true cognitive science discipline.

Once upon a time, philosophy was characterized as the queen of the sciences. But then various special sciences (psychology, linguistics, etc.) emerged to address specific phenomena, leaving philosophy with the remnants for which no special science had developed. Philosophers still talk about mind, language, or computation, but often these discussions seem to bear little relevance to the investigators in the autonomous disciplines. Indeed, for those discussions that do not afford linkage to the rest of cognitive science, the philosophical banter might best be ignored by the major participants of cognitive science. A favorite tactic of philosophers is the construction of arcane thought experiments designed to reveal supposedly necessary conceptual connections. One classic example posits Swampman, an exact replica of you who somehow emerged out of swamp gas a millisecond ago as you were vaporized (Davidson, 1987). Does Swampman, equivalent to you molecule for molecule, have the same thoughts as you? Presuming that a negative answer is obvious, Davidson construes the thought experiment as demonstrating that thoughts require a history of engagement with an environment. Others, with different intuitions, reach a different conclusion. Underlying this practice of arguing for substantive conclusions from thought experiments is the claim that our concepts, such as the concept of ''thought,'' are fixed and the intuitions we appeal to in evaluating thought experiments reveal the true nature of concepts. In the context of cognitive science, with its ingenious research programs challenging the notion of fixed concepts, the ensuing philosophical arguments often seem beside the point.

There is, however, a different approach to philosophy, characterized by philosophers themselves as naturalistic. It offers conclusions that have a direct bearing on cognitive science and its practitioners can profit from closer engagement with the rest of cognitive science. Although it has roots much earlier in the history of philosophy, its modern incarnation began when Quine (1969) proposed that naturalized epistemology could situate itself within psychology and yet still address fundamental questions about how knowledge is acquired. (That would include knowledge obtained by psychology, which in his era was behaviorist in orientation.) Extending Quine's maneuver to other areas of philosophy, a naturalized philosophy of mind would become the philosophy of cognitive science and draw upon results in cognitive science as well as its own inquiries to probe the nature of mental phenomena. Likewise, naturalized philosophy of science would study actual lines of

research in the cognitive (and other) sciences to develop accounts of how science itself works. These accounts in turn may support normative prescriptions for further research in cognitive science. One might recognize philosophical discussions conducted in a naturalistic manner not as arcane intellectual exercises, but as theoretical and methodological contributions to cognitive science. As a bonus, unlike other sciences subjected to philosophical inquiry, cognitive science offers a two-way street: Its theoretical frameworks and findings are a resource philosophers can draw upon to add nuance and techniques to their own subsequent inquiries.

In some domains, philosophers have already engaged in this two-way engagement with cognitive science. Thagard, for example, has collaborated with psychologists and AI researchers in the investigation of induction (Holland, Holyoak, Nisbett, & Thagard, 1986) and drawn on cognitive science tools to understand conceptual change (Thagard, 1992). Nersessian (2008) shows how cognitive science research on mental models can be extended to understand the development of scientific models and the creation of new scientific concepts. Giere (2006) has applied cognitive science analyses of distributed cognition to large-scale physics projects, extending thereby both the understanding of distributed cognition and of scientific reasoning. In what follows I briefly consider three topics that have tended to be pursued primarily by philosophers (two in philosophy of mind, one in philosophy of science) that bear on cognitive science and discuss how they might become more fully integrated with cognitive science in the future.

2. Philosophy of mind: The mind–body problem

Philosophers have long debated how the mind relates to the physical body (and how mind itself is to be construed). In recent decades, the accounts offered have closely paralleled developments in the cognitive and brain sciences. For example, pioneering researchers in the early decades of the 20th century developed techniques for recording from neurons and nerve fibers. This resulted in the identification of different types of nerve fibers and the recognition that each transmits a different sort of information (e.g., C-fibers carry pain sensations). This was part of the inspiration for several philosophers (e.g., Smart, 1959) who advanced the identity theory—the claim that mental processes are identical with brain processes. (Its mantra was ''pain is C-fiber firing.'') The brief ascendancy of identity theory coincided with the period in which psychologists, linguists, and the founders of AI were developing mentalistic explanations of cognitive activity. Although some of these pioneers sought to incorporate what was known about the brain in developing their accounts, they were limited by the paucity of tools for linking human mental activities to the brain. (This was a major reason the neurosciences ended up playing such a minor role in the first decades of cognitive science.) Cognitive approaches inspired functionalism as an alternative to identity theory in philosophy. Functionalism's key claim is that interactions between mental states—rather than their constitution—determine their identity. Within philosophy, functionalism was widely regarded as offering the correct account of mind–body relations for cognitive science. However, functionalism itself confronted problems, such as an

apparent inability to account for the qualitative nature of experience. It was claimed that the qualitative character of the experience of seeing red could be switched with the experience of seeing blue but leave functional relations the same as before the switch. This led a few philosophers to repudiate functionalism (and some, e.g., Chalmers, 1996, even to despair of the possibility of any physicalist account of mind).

Putnam (1967) introduced the best-known philosophical argument for choosing functionalism over the identity theory. Just as computer programs can be compiled to run on different hardware, he argued that mental processes such as hunger and pain can be instantiated (realized) in all sorts of physical systems, including other species and computers, and so could not be identified with any one realization. In generalizing this multiple realizability objection to all special sciences, Fodor (1975) presented money as a compelling example of a type of entity that is the topic of many important generalizations but can be realized by an unlimited range of physical entities. Through the 1980s and 1990s, the multiple realizability argument seemed so intuitively obvious that it was seldom challenged (for an early challenge, see Thagard, 1986, who appealed to parallel computing to illustrate the differences the physical realization may make to the processing a system can perform). The claim that the same mental processes can be realized in very different ways in different species conflicts with practices in neuroscience, however, which often assume a comparative perspective. In delineating brain regions, for example, neuroanatomists since Brodmann have developed common maps for a broad range of species. Neurophysiologists have developed accounts of mammalian visual processing by conducting research primarily on cats and monkeys. They have been untroubled by what differences exist across species in brain regions and processes, as their similarities were more than sufficient to support a common system of categories. Interestingly, though, the functionalist is similarly grouping mental states when he or she treats a given mental process in different organisms as the same despite clear behavioral differences, yet at the same time she refuses to make the same move with respect to the underlying neural processes. The alleged examples of multiple realization of mental processes arise only when one lumps mental processes but splits brain processes (Bechtel & Mundale, 1999). In the last decade, several philosophers have begun to challenge the multiple realizability argument and re-explore the merits of the identity theory.

The philosophical debate between identity theory and functionalism has the potential to inform cognitive science as it confronts the issue of how to situate itself vis-a-vis the recent development of cognitive neuroscience. The methodological obstacles to relating cognitive processes to brain processes faced at the inception of cognitive science were partly overcome with the introduction of technologies for functional neuroimaging beginning in the 1980s. This has resulted in a split within cognitive science that closely mirrors the divide between identity theorists and functionalists in philosophy of mind. Like identity theorists, those who embrace the techniques of cognitive neuroscience identify cognitive mechanisms with brain mechanisms, whereas those who, like philosophical functionalists, advocate that cognitive science remains autonomous from the brain sciences by developing and testing their accounts of cognitive processes while making no appeals to the brain. For some this amounts to benign neglect, whereas others regard themselves as defending cognitive science from a neural imperialism that threatens to supplant cognitive inquiry with brain studies.

One way to advance this discussion in cognitive science is to develop in greater depth the corresponding philosophical accounts. For example, certain philosophers (Armstrong, 1968; Lewis, 1980) have long embraced the notion of integrating functional considerations with the identity theory; they do so by treating brain states as instantiating mental states characterized functionally (e.g., in terms of their interrelatedness). This approach can be further advanced by recognizing that the neurosciences are themselves concerned with functional relations between brain states. As developed in Bechtel and McCauley's (1999) heuristic identity theory, identity claims can then be viewed as a heuristic strategy for advancing both cognitive and neuroscience inquiries by allowing information developed in either to guide and constrain investigations in the other. As I will discuss briefly in Section 3, when the goal is viewed as developing accounts of mental mechanisms, researchers may employ the resources of both cognitive and neural inquiries without subordinating one to the other.

3. Philosophy of mind: The nature of representations

Contemporary philosophical interest in mental representations, although in part elicited by the invocation of representations in cognitive science models, is historically rooted in Brentano's contention that intentionality is the mark of the mental. Attention to representation in cognitive science has focused primarily on what philosophers characterize as the *vehicles* of representation—here, the mental states that serve a representational function. *Intentionality*, on the contrary, refers to the relation between a vehicle and its content. This surprisingly slippery construct has numerous interpretations; most relevant here is the idea that mental representations typically have as their content objects or events outside the mind. (One much-discussed complication is that the connection to content cannot be an ordinary relation in that the content—unicorns being a commonly cited example—need not exist.) Philosophical treatments of mental representation have considered both vehicle and content, but it is primarily the debates about vehicles that have had impact on cognitive science to date.

In his pioneering *Language of Thought*, Fodor (1975) argued that cognitive science explanations are committed to the assumption that cognitive agents possess an innate system of language-like representational vehicles (a language of thought). Cognitive science's commitments must go well beyond the mere claim that mental representations exist (he labels this claim the *representational theory of mind*), a view that can be embraced even by associationists who posit associations between ideas (i.e., representations). Cognitive science must also, he contends, be committed to causal (computational) processes operating on these representational vehicles (he calls this the *computational theory of mind*). For such operations to produce thought, Fodor continues, the vehicles must combine into subject–predicate structures, and the predicate that is the vehicle for a particular property must be distinct from other vehicles (i.e., have its own ''shape'' or form). Then rules that apply only to those forms, not to content, can enable inferences about specific properties. Fodor also contends cognitive science must accept that such a language-like representational system (mentalese) is innate. His argument begins with the assumption that the only plausible account of

learning (including the learning of natural languages) requires hypothesis testing. To formulate a hypothesis about a new concept to be learned, a cognitive agent must already possess or represent the concept; hence, it is innate and need not be learned after all. (Fodor also contends that perception requires formulating hypotheses about what is being perceived and testing these against sensory input. Hence, perception too requires prior possession of representations.)

Initially, nativism was regarded as the most controversial aspect of Fodor's language of thought hypothesis. The computational theory of mind seemed less problematic as the major AI models all invoked formal operations over representational vehicles. The reemergence of connectionist models in the 1980s, however, posed a challenge—although patterns of activation over units in networks were construed as representational, these representations were not structures that were operated on by formal rules. Fodor and Pylyshyn (1988) argued that this was a fatal flaw. They allowed that connectionist models might figure in explaining the realization (implementation) of appropriately structured representations, but that mental activity involved the realization not just of these structures but also of operations defined over them. In the course of confronting connectionism, Fodor and Pylyshyn advanced new arguments for positing language-like representations. They argued that thought, like language, is both productive (it is always possible to think new thoughts) and systematic (any agent capable of thinking one thought can think other thoughts that are appropriately related to the first). Only representational vehicles with compositional syntax, they maintained, could support these features of thought.

Philosophers and other cognitive scientists who were attracted to connectionism have questioned whether thought is as productive and systematic as Fodor and Pylyshyn claimed and contended that systems without language-like vehicles could account for the degree of productivity and systematicity that is exhibited. The ability of connectionist networks to learn in ways other than by testing hypotheses has also resulted in challenges to Fodor's arguments that such representations must be innate (for discussion, see Bechtel & Abrahamsen, 2002, Chapter 6). As a result, although the language of thought hypothesis is still embraced by many philosophers and cognitive scientists, it is no longer regarded as the ''only game in town,'' as Fodor maintained.

Whereas philosophical claims about vehicles of representation have engaged cognitive scientists, accounting for their content (their connection to entities and events they represent) has been a more exclusively philosophical pursuit. Again, Fodor has been a central contributor. On the one hand, he claims that in embracing the computational theory of mind, cognitive science must ignore content—a tactic Putnam called *methodological solipsism*. This follows from the fact that formally characterized computational processes operated solely on vehicles, not their content (Fodor, 1980). On the other hand, Fodor holds that it is important to explain how it is that representations do have content. To explain this, Fodor was initially attracted to a causal account whereby the content caused the occurrence of the vehicle in the mind (an approach forcefully advanced by Dretske, 1981, who viewed mental states as carrying information about the events that caused them). Yet an important constraint on any account of content is that it must be able to account for misrepresentation: We can misrepresent a stick as a snake and can represent things that do not actually exist. This

poses a challenge for causal accounts. Even in cases of supposed misrepresentation, something triggers the representation, and if we count that as the content, as the causal theory proposes, there is no misrepresentation. A common strategy for surmounting this problem has been to privilege some causes as providing the content and then to allow for misrepresentation when other causes trigger the representation. Millikan (1984) proposed that natural selection provides a means to specify the appropriate causal linkage—the content of a representation was what it was selected to represent (see Fodor, 1987, for arguments against this approach).

Although these philosophical debates about content have attracted interest among other cognitive scientists only somewhat episodically, the issue is an important one for cognitive science. The importance and difficulty of this challenge is illustrated by Dreyfus's (1992) claims about what computers cannot do and by Searle's (1980) well-known Chinese Room argument. Both contend that accounts of mind that treat it as a manipulator of representational vehicles cannot capture the meaning of thoughts. Dreyfus in particular, drawing upon Heidegger, suggests that a focus on cultivated skills and our social nexus offers a more promising avenue to explaining human mental practices. On such a view, representing might first be viewed as a culturally supported practice involving public symbols that we learn to use privately only later. Philosophers (e.g., Wheeler, 2005) are pursuing this approach and in the process engaging with recent cognitive science research on embodied and situated cognition.

Another line of criticism of standard approaches to representation is found among advocates of dynamical systems theory (DST), who argue that computational approaches offer too static a view of the processes occurring within the brain (van Gelder & Port, 1995). From the DST perspective, dynamical internal processes are complexly coupled with the world with which the organism is interacting (or mistakenly taking itself to be interacting). This suggests a different perspective on representation that does not encounter the problems faced by accounts that begin with formal symbols and must separately link them back to the world. A related promising approach to linking representations with their contents involves viewing the mind–brain as consisting of control systems whose internal processes (representations) provide organisms with means of regulating their bodies in relation to things and events in the world (for suggestions along these lines, see Bechtel, 2008, 2009; Grush, 2004). Although these approaches construe representations rather differently than traditional cognitive science accounts, they afford productive engagement with research endeavors in cognitive science emphasizing the sensory and motor grounding of concepts and cognitive processes (Barsalou, 2008).

4. Philosophy of science: Explanations

Philosophy of science concerns itself with how science itself operates to generate knowledge. A central question is what counts as a scientific explanation. An influential tradition in philosophy of science that developed in the early decades of the 20th century appealed to logic and linguistic representations to characterize explanations: They involved statements

of laws from which descriptions of events to be explained could be formally derived. The resulting D–N (deductive–nomological) account was presented as normative—specifying what any science should aim for. A naturalistic approach to philosophy of science, in contrast, turns first to the practices of science itself. Any norms to be advanced must evolve from accurate accounts of science. In the cognitive sciences, or the life sciences more generally, one finds very few appeals to laws, and such laws as are advanced (or effects, as they are often called in psychology) merely characterize the phenomenon to be explained (Cummins, 2000). When cognitive scientists present explanations, they often appeal to mechanisms that are thought to produce the phenomenon—they do not simply characterize it.

Recently, a number of naturalistically inclined philosophers of science have sought to characterize mechanistic explanations and how they figure in scientific practice. Although terminology differs, these accounts agree that mechanisms are taken to consist of parts that carry out different operations and are organized so as to produce the phenomenon of interest (Bechtel & Abrahamsen, 2005). In reasoning about mechanisms, investigators not only sometimes employ linguistic representations and arguments but may also utilize diagrams and mental simulations. The questions of how diagrams represent and how people simulate processes in the world are research topics within cognitive science, and philosophers seeking to understand mechanistic explanations avail themselves of this research. An important characteristic of mechanistic research in the life sciences is that usually it is concentrated on model systems (e.g., *Drosophila*), from which researchers may later attempt to generalize to other systems. But model systems are commonly chosen not for their typicality but because they afford investigation. Generalization does not take the form of treating as universal what is found in a particular case (as simplistic accounts of induction would have it), but of identifying similarities and differences between mechanisms. In biology, phylogenetic descent relations result in mechanisms being conserved, yet exhibiting considerable variation across species.

The development of mechanistic accounts of specific phenomena often requires recruiting a variety of lines of research, which may rely on different methodologies and instruments. This may involve interdisciplinary engagement, such as is found in cognitive science as well as many other domains of modern science. More traditional accounts of integration in science emphasized theory reduction—the derivation of one science's laws from those of another science—and one still finds in philosophy frequent discussion of whether, for example, psychological laws reduce to those of neuroscience. Reduction from a mechanistic perspective, in contrast, involves securing information about the parts and operations that enable a mechanism to produce a phenomenon. But understanding the mechanism also requires understanding how the parts and operations work together and how the mechanism as a whole relates to its environment. If the organization is complex and operations are orchestrated in time, tools of complexity theory and DST may need to be brought to bear in dynamic mechanistic explanations (Bechtel & Abrahamsen, in press). Moreover, the way the mechanism is situated in its environment (e.g., when a memory mechanism is operating in conditions favoring formation of false memories) can be just as important to a full understanding of a phenomenon as operations within the mechanism. Causal processes occur at different levels and each contributes to the understanding of how mechanisms function.

Naturalistic philosophy of science has potentially rich connections to cognitive science. First, cognitive science inquiries, like those in other sciences, serve as the subject matter for analyses in philosophy of science. Second, philosophers can draw upon results of cognitive science in characterizing the cognitive activities of scientists. Third, to the degree that philosophers of science can elicit, from their examinations of actual science, norms about the forms explanation may take and the modes of inquiry likely to generate successful explanation, they may be able to offer productive advice to cognitive scientists as they evaluate proposed explanations and researcher strategies. (See Bechtel, 2009, for further development of issues raised in this section.)

5. Conclusions

I have briefly described recent philosophical contributions to two issues in philosophy of mind (the mind–body problem and the nature of mental representations) and one in philosophy of science (explanation). All three are of potential importance to cognitive science, although they have often been discussed in isolation from the other cognitive sciences, rendering philosophy an outlier. In each case I have suggested ways in which naturalized approaches to these issues can integrate philosophy more fully into cognitive science. Stances on the mind–body problem can be important as cognitive science negotiates its relation to cognitive neuroscience. Philosophers have long been engaged in evaluating potential vehicles of mental representation, but new approaches to content suggest how to understand the relation of cognitive agents to their environments. Finally, the recently formulated accounts of mechanistic explanation both characterize the explanatory project of cognitive science and suggest ways in which multiple disciplinary perspectives in cognitive science may be integrated in studying a common mechanism.

Acknowledgments

References

Armstrong, D. M. (1968). *A materialist theory of mind*. London: Routledge and Kegan Paul.
Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, *59*(1), 617–645.
Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. London: Routledge.
Bechtel, W. (2009). Constructing a philosophy of science of cognitive science. *Topics in Cognitive Science*, *1*(3), 548–569.
Bechtel, W., & Abrahamsen, A. (2002). *Connectionism and the mind: Parallel processing, dynamics, and evolution in networks* (2nd ed.) Oxford, England: Blackwell.

Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences*, *36*, 421–441.

Bechtel, W., & Abrahamsen, A. (in press). Complex biological mechanisms: Cyclic, oscillatory, and autonomous. In C. A. Hooker (Ed.), *Philosophy of complex systems. Handbook of the philosophy of science* (Vol. 10). New York: Elsevier.

Bechtel, W., & McCauley, R. N. (1999). Heuristic identity theory (or back to the future): The mind-body problem against the background of research strategies in cognitive neuroscience. In M. Hahn & S. C. Stoness (Eds.), *Proceedings of the 21st Annual Meeting of the Cognitive Science Society* (pp. 67–72). Mahwah, NJ: Lawrence Erlbaum Associates.

Bechtel, W., & Mundale, J. (1999). Multiple realizability revisited: Linking cognitive and neural states. *Philosophy of Science*, *66*, 175–207.

Chalmers, D. (1996). *The conscious mind*. Oxford, England: Oxford University Press.

Cummins, R. (2000). ''How does it work?'' versus ''what are the laws?'': Two conceptions of psychological explanation. In F. Keil & R. Wilson (Eds.), *Explanation and cognition* (pp. 117–144). Cambridge, MA: MIT Press.

Davidson, D. (1987). Knowing one's own mind. *Proceedings and Addresses of the American Philosophical Association*, *60*, 441–458.

Dretske, F. I. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press∕Bradford Books.

Dreyfus, H. L. (1992). *What computers still can't do: A critique of artificial reason*. Cambridge, MA: MIT Press.

Fodor, J. A. (1975). *The language of thought*. New York: Crowell.

Fodor, J. A. (1980). Methodological solipsism considered as a research strategy in cognitive psychology. *The Behavioral and Brain Sciences*, *3*, 63–109.

Fodor, J. A. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*. Cambridge, MA: MIT Press.

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*, 3–71.

van Gelder, T., & Port, R. (1995). It's about time: An overview of the dynamical approach to cognition. In R. Port & T. van Gelder (Eds.), *It's about time* (pp. 1–43). Cambridge, MA: MIT Press.

Giere, R. G. (2006). *Scientific perspectivism*. Chicago: University of Chicago Press.

Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, *27*, 377–396.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). *Induction: Processes of inference, learning and discovery*. Cambridge, MA: MIT.

Lewis, D. (1980). Mad pain and Martian pain. In N. Block (Ed.), *Readings in philosophy of psychology* (Vol. 1, pp. 216–222). Cambridge, MA: Harvard University Press.

Millikan, R. G. (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.

Nersessian, N. (2008). *Creating scientific concepts*. Cambridge, MA: MIT Press.

Putnam, H. (1967). Psychological predicates. In W. H. Capitan & D. D. Merrill (Eds.), *Art, mind and religion* (pp. 37–48). Pittsburgh: University of Pittsburgh Press.

Quine, W. V. O. (1969). Epistemology naturalized. In W. V. O. Quine (Ed.), *Ontological relativity and other essays* (pp. 37–48). New York: Columbia University Press.

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, *3*, 417–424.

Smart, J. J. C. (1959). Sensations and brain processes. *Philosophical Review*, *68*, 141–156.

Thagard, P. (1986). Parallel computation and the mind-body problem. *Cognitive Science*, *10*(3), 301–318.

Thagard, P. (1992). *Conceptual revolutions*. Princeton, NJ: Princeton University Press.

Wheeler, M. (2005). *Reconstructing the cognitive world: The next step*. Cambridge, MA: MIT Press.