To appear in S. Sarkar (ed.)
***The Philosophy of Science: An Encyclopedia*.**
New York: Routledge Press, 2005.

PHILOSOPHY OF PSYCHOLOGY

Robert A. Wilson
University of Alberta

1. Introduction

In the good old days, when general philosophy of science ruled the Earth, a simple division was often invoked to talk about philosophical issues specific to particular kinds of science: that between the natural sciences and the social sciences. Over the last 20 years, philosophical studies shaped around this dichotomy have given way to those organized by more fine-grained categories, corresponding to specific disciplines, as the literatures on the philosophy of physics, biology, economics and psychology--to take the most prominent four examples--have blossomed. In general terms, work in each of these areas has become increasingly enmeshed with that in the corresponding science itself, increasingly naturalistic (in at least one sense of that term), and in my view, increasingly interesting.

The philosophy of psychology is concerned with mind and cognition. When psychology cut itself loose institutionally and professionally from philosophy in the late nineteenth- and early twentieth-centuries, it was the discipline that predominantly studied mind and cognition. This has changed over the last 30 years. With the development of artificial intelligence, cognitive anthropology, linguistics, and neuroscience--perhaps, together with psychology, best referred to collectively as "the cognitive sciences" (Wilson and Keil 1999)--philosophers of psychology have found themselves both drawing on and contributing to scientific work in this more interdisciplinary milieu. There are two consequences of this. The first is that the field has become increasingly entwined with the philosophical aspects to cognitive science. My own view is that we do greater justice to the interdisciplinary motivations behind cognitive science by placing an emphasis on the cognitive sciences, rather than on any foundational assumptions that constitute a paradigm, cognitive science (see COGNITIVE SCIENCE). Thus, I view the philosophical aspects to the cognitive sciences as occupying the greater part of the philosophy of

psychology (cf. Wilson 1999). The second is that the more lively areas or topics of contemporary discussion in the philosophy of psychology are quite diverse, including (for example) philosophical issues in neuroscience, the nature and physical bases of consciousness, the evolution of mind, and the ontogenetic and phylogenetic development of intentional states in human agents.

Despite the first of these points, and contributing to the second, the material that philosophers of psychology discuss also covers questions about the mind and areas of psychology that even my own pluralistic (licentious?) conception of the cognitive sciences excludes. Here I have in mind debates over the scientific status of psychoanalysis, questions about the foundations of the taxonomy of psychopathology, and discussions of the nature of social psychology, all of which concern areas of psychology other than cognitive psychology.

What further complicates any simple characterization of work in the philosophy of psychology, and to some extent what distinguishes it from the other "philosophy of x" studies within the philosophy of science, is its close relationship to a traditional area of philosophy, the philosophy of mind, that has not typically viewed itself as a part of the philosophy of science at all. Thus, many of the topics that philosophers of psychology discuss that arise from their reflection on the cognitive sciences have analogues in traditional philosophy of mind. For example, concerns about the causal role of semantic or representational level properties in computational theories of cognition echo the more general problem of mental causation; many of the issues about the nature of cognitive architecture that separate, for example, "classic" from connectionist approaches to cognitive architecture, are also reflected in the historical debates between rationalists and empiricists. Perhaps because the nature of the mind has been one of the central issues in metaphysics and epistemology throughout the history of philosophy, the connections between the philosophy of psychology and philosophy more generally are more extensive than in any other disciplinarily specialized area of the philosophy of science.

In what follows I shall attempt both to convey something of the flavor of three topics within the philosophy of psychology that have dominated the field over the last 20 years: intentionality, cognitive architecture, and

consciousness. I will also briefly discuss another pair of more specific topic clusters that represent novel and perhaps trend-setting topics for future research.

2. Intentionality and Mental Representation

The postulation of mental representation has been central to the cognitive sciences throughout their history, and it is easy to see why. Human agents do not simply or reflexively respond to their environments, but are equipped with some internal, mediating mental machinery, machinery that is sensitive to what is in the environment but which has enough complexity to it to thwart any attempt, (e.g., that made by behaviorists), to exhaustively characterize it in terms of that environment (e.g., in terms of stimulus-response pairs). Mental representations play precisely such a mediating role, both containing information about the world, and combining to guide an individual's behavior in that world (see INTENTIONALITY).

The form that mental representation takes in our common sense, folk psychology is *propositional*: agents have beliefs and desires, where each of these mental states can be thought of as an attitude to a proposition. Since psychology has built on such folk psychological representations since its inception--from the Freudian extension of folk psychology from conscious to unconscious states, to work on stereotypes and schemata in social psychology, to classic AI models of human problem solving or reasoning--propositional representation has been a focus of discussion within the philosophy of psychology. In fact, due to its prominence, many general discussions of mental representation have been cast exclusively in terms of propositional representation, or even its folk psychological guise. Here are three of the central issues in the literature, and a sampling of positions that have been adopted with respect to those issues:

How many kinds of mental representation are there? Much of the debate over mental imagery (Kosslyn 1994) has focussed on the reality of mental images and their relationship to propositional representations. There has also been more recent discussion of the extent to which mental representations are "local" as opposed to "distributed" in their nature (Haugeland 1991). The role of language in mental representation, and thus thought, has also structured a range of related debates, such as that over the language of thought

hypothesis, and the question of the form that mental representation takes in non-linguistic creatures, such as human infants and non-human animals.

What determines a representation's content? Three chief answers to this question have been entertained: conceptual role or procedural semantics, causal or informational theories, and teleological theories (see also COGNITIVE SCIENCE). The first of these is typically internalist in that mental content is determined entirely by intrinsic, physical properties of the agent or system. But the most pervasive views here are externalist, and as such they have reinforced externalist views of psychology and psychological explanation (Wilson 1995, 2002). Both causal and teleological views allow an individual's historical, social and physical location to partially determine what content its representations have. An alternative form of externalism that departs from the sort of realism about mental representation that has been taken for granted by the three chief views here is a conventionalism about the nature of representational content, defended in different ways by Dennett (1987) and Horst (1996).

Is mental representation dispensable within the cognitive sciences? Stephen Stich (1983) was an early defender of the view that the cognitive sciences could be (indeed, should be) content-free; the Churchlands, through various publications, have expressed an alternative, neuroscientifically-inspired form of eliminativism about mental representation. Both of these forms of eliminativism about mental representation have pitched their critiques at the sorts of representations posited by folk psychology. Proponents of connectionist architectures and, more recently, of dynamic approaches to cognition have also often introduced their views as avoiding the postulation of mental representation. But as the names "distributed" and "dynamic" representations suggest, such approaches do not necessarily imply the rejection of all forms of mental representation, and the place of mental representation within them remains a topic of continuing interest (Érdi 2000).

3. Cognitive Architecture and Processing

If debates over the nature of mental representation concern *what* it is that cognition ranges over, those over cognitive architecture and processing concern *how* it is that cognition proceeds. Part and parcel of the "cognitive revolution" of the late 1950s that formed the basis for the cognitive sciences was the conceptualization of cognitive processing as a form of *computation*.

This view, computationalism, has received both general and somewhat vague characterizations ("cognition is computation"), as well as more specific formulations ("cognition is explicit symbol manipulation") that are tied to particular research programs, the best known of which is associated with Allen Newell and Herb Simon as the *physical symbol system hypothesis*: "a physical symbol system has the necessary and sufficient means for general intelligent action". Central to any account of cognitive processing is a commitment to the nature of the basic design of the cognitive system, the *cognitive architecture* of that system, and hypotheses about cognitive architecture have usually been formulated as explicit computational models that generate behavior that approximates some aspect of (often human) cognitive behavior. On Newell and Simon's own view, production systems, which consist of chains of condition-action rules defined over data structures, form the heart of human cognitive architecture, and the types of behaviors to which their computational models were applied most extensively were problem solving and reasoning (Newell and Simon 1972). Variations on this general view were predominant in much of AI and psychology until the 1980s, and the philosopher perhaps most firmly associated with this sort of "rules and representations" approach to cognitive architecture is Jerry Fodor (1981, 1987; see also COMPUTATIONAL THEORIES OF MIND).

Over the past 20 years, connectionism has come to represent a general alternative to the rules and representation approach. The basic idea of connectionist architectures and the neural network models that correspond to them is that cognition involves the adjustment of weighted connections between many relatively simple processing units through a process of feedback from environmental inputs (learning). Although these basic units are often compared to neurons, the bulk of the psychological work to which philosophers appeal (e.g., on modeling the acquisition of the past tense in English) involves processing units that are on the wrong scale to be very neuron-like (see CONNECTIONISM).

The most fruitful work within the computational paradigm, broadly construed, involves models that appeal to aspects of both rules and representations and connectionist architectures. A common suggestion is that the former handles "higher" cognitive functions, such as problem solving, while the latter are applicable to "lower level" cognition, such as pattern recognition. But more truly integrative models of cognitive architecture focus

on the role that *probability* has within computational models; for example, Boltzmann machines, developed within the neural network paradigm, are essentially identical to Bayesian networks developed within traditional artificial intelligence (Jordan and Russell 1999, Pearl 2000).

Dynamic approaches to cognition attempt to pose a more radical challenge to these two views of cognitive architecture and their corresponding paradigms for the cognitive sciences. The chief idea of dynamicism is that cognitive systems are a form of dynamic system, a system that exists in real time and whose movement over time is not governed by any special computational principles (Port and van Gelder 1995). On the dynamic conception of cognitive processing, internalized rules and symbols do not play any special role in cognition; rather, cognition proceeds through the settling of the cognitive system into an equilibrium state. The mathematical equations that govern such processes are not internalized within the cognizer any more than Newton's laws of motion are internalized in the object's whose behavior they govern. The dynamic approach has thus challenged both the representational and computational dimensions to standard cognitive science, and it also suggests that cognitive systems are fundamentally *embedded* or *embodied*, a point to which we will return.

The development of connectionist architectures has led many philosophers of psychology to rethink a range of issues concerning the nature of cognitive processing. Many of these concern the nature of mental representation, as we have already seen, but the rise of connectionism has also generated more general discussions, such as those over the nature of computation (including the relationship between computational models and computation) and the role that cognitive neuroscience has to play in addressing some of these questions about large-scale cognitive organization. Despite the fact that most of the neural network models of influence within the cognitive sciences are not neurally very realistic, connectionist architectures have redirected attention to the brain itself, particularly as non-invasive techniques of imaging, such as positron emission tomography (PET) and functional magnetic resonance imagine (fMRI), have allowed researchers to explore the activity of the brain in real time.

One pair of related themes concerning cognitive architecture has been the modularity of cognitive design and the localization of mental processes.

Fodor (1983), crystallizing and generalizing a view of the mind articulated by Noam Chomsky as a part of his approach to generative grammar in linguistics, rekindled interest in a modular view of cognitive capacities of the sort introduced originally by Franz Gall almost two hundred years earlier. According to Fodor's view, many such capacities were *domain-specific* and *encapsulated*: roughly, cognition is structured so that particular mental organs are sensitive only to specific kinds of inputs, and are insulated from the causal influence of the operation of other mental organs. Fodor's own view here was that such a view of the mind held only of input systems--the five senses, plus language, according to Fodor--and he cautioned against the extension of the view to "central systems". This caution has been largely ignored as developmental psychologist have postulated Fodorean modules for the domains of physics, number, biology, and psychology, and evolutionary psychologists have endorsed what has become known as the *massive modularity* thesis, the claim that the mind is overwhelmingly modular, with the number of modules running into the hundreds if not thousands (see EVOLUTIONARY PSYCHOLOGY). Philosophers have had much to say about these topics, particularly about the "theory of mind" within developmental psychology and evolutionary psychology in general.

It has typically been assumed that modules were *physically localized* in the brain, roughly in the way in which other bodily organs, such as the heart or the kidney, were so localized. As Fodor himself pointed out in his brief discussion of the "fixed neural architecture" (1983:98-99) associated with modules, we might articulate this assumption in terms of broader systems that are somewhat distributed throughout the brain. But the basic idea is that functionally individuated modules have neural hardware specifically dedicated to the function that they perform. As PET and fMRI have been increasingly used in experimental investigations of cognition, data on such localization assumptions has accumulated, though it is worth mentioning that these methods themselves have often been used in ways that presuppose a basically localistic view of cognitive function (Uttal 2001). Dan Lloyd (2000) has presented a striking, even if preliminary, meta-analysis of the data across independent studies, arguing that this data supports the claim that the brain is a distributor processor, and refutes the stronger, "localistic", modularity hypotheses common in the field. See also NEUROBIOLOGY.

4. Consciousness

Consciousness has been a buzz topic in the philosophy of psychology for the past 10 years, returning to occupy center stage after a long absence, and commanding the attention both of philosophers of science (i.e., of psychology) and traditional philosophers of mind. Amongst the latter, there has been an unblushingly (indeed, celebratory) *a prioristic* strand, with a focus on the challenge that consciousness, phenomenal states, and qualia pose to views, such as physicalism and functionalism, that continue to operate as working assumptions for many within the cognitive sciences. A work that has galvanized such discussion is David Chalmers' *The Conscious Mind* (1996), a book whose central conclusions echo the skepticism about physicalism associated with well-known, earlier papers by Thomas Nagel and Frank Jackson, and whose emphasis on conceivability arguments and what they putatively show about the limits to the scientific study of consciousness has fuelled some interesting debate over the role of conceptual analysis within a naturalistic account of the mind. Philosophical work on consciousness has also reacted to attempts to "quine qualia", developed representationalist views of the nature of phenomenal experience, and debated the idea, due to David Armstrong and David Rosenthal, that consciousness just is awareness, so that a conscious mental state is some sort of second- (or, in general, higher-) order mental state. A succinct overview of this work is provided by Levine (1997).

Within the cognitive sciences themselves during roughly the same period consciousness has also had a renaissance, with much of this literature focussed on the phenomena of visual awareness and attention. It has speculated about the function(s) and evolutionary origins of consciousness, as well as utilized a variety of neural techniques to try to pinpoint the parts of the brain that are most directly causally responsible for conscious experience, In an influential paper in 1990, Christof Koch and Francis Crick advocated that the time was ripe for neural speculations about consciousness, and proposed (building on the work of von Malsburg and others) that 40 hertz oscillatory cycles in the brain, particularly in the visual cortex, were especially important to consciousness, A detailed, recent empirical account of consciousness has been offered by Rodney Cotterill (1998), which emphasizes the relationship between consciousness and movement and the importance of *timing* to consciousness. Cotterill offers an integrated psychological and neurological view of the bases for consciousness that posits a triangular neural circuit linking the posterior lobes, the premotor area of the frontal lobes, and the

nucleus reticularis thalami between the thalamus and the medulla oblongata as the neural basis for conscious experience.

One obvious question to be asked concerns the relationship between such work on consciousness and that on the nonconscious, representational mental states that have been studied within the bulk of the cognitive sciences over the preceding 30 years. *Representationalism* about conscious states constitutes one sort of answer, for it holds that qualitative states just *are* representational states. Indeed, one of the motivations for representationalism is to deflate the commitments that one makes in admitting conscious mental states as well as intentional states to one's ontology. Another type of answer is provided by John Searle (1992), who has defended what he calls *the connection principle*, which says that unconscious mental states must be, in principle, accessible to consciousness. This principle has one of two implications for traditional cognitive science: either the states it posits do not exist, or those states are, contrary to what most of those investigating them believe, accessible to consciousness.

Given the disparate writings on consciousness, it is no surprise that some of these have become more explicitly self-reflective. Perhaps the best-known of these is Ned Block's "On a Confusion About a Function of Consciousness" (1995), which introduced the distinction between what Block calls phenomenal consciousness--the what-it's-likeness of mental experience--and access-consciousness, the feature of mental experience that allows its reportability. One suggestion is that mental states such as pain and sensations are p-conscious, while those such as occurrent thoughts are a-conscious; another is that the former is really the subject of the Nagel-Jackson-Chalmers inspired literature, while the latter is what cognitive scientists investigate. Block himself introduced the distinction to critique claims, especially in the psychological literature, that were often made about the function of phenomenal consciousness that relied implicitly only on data about access-consciousness. Philosophers remain divided over whether Block's distinction makes sense of much of the consciousness literature, or whether it constitutes a confusion about consciousness itself. See also CONSCIOUSNESS.

5. Pain, Psychopathology, and Color

One of the concomitant products of the extended consciousness-fest has been work on topics concerning particular phenomenal states. *Color* is perhaps the

most richly-mined of these, beginning with C.L. Hardin's *Color for Philosophers* (1988), which significantly raised the bar regarding the level of empirical detail relevant to philosophical discussions of color. From its characterization as a secondary quality in 17th-century corpuscularian philosophy and science, color has constituted both an epistemic and an ontological puzzle: just what is color, in the world, and does our epistemic access to it constitute some sort of privileged knowledge? Some of the recent work on color processing in the cognitive sciences suggests that color constitutes at least as much an enigma to accounts of cognitive processing. For example, it now appears that there is no place or system in the central nervous system that is modularly dedicated to process color, and this has led some philosophers to rethink the evolutionary function of color perception and its role within the perceptual life of the individual (cf. Matthen 1999).

While psychopathology itself is not a new topic for philosophers, work here has taken a novel turn as a by-product of the focus on consciousness. Conscious experience sometimes deviates from its normal course. Philosophical issues abound here, whether it be in cases of blindsight in patients with severed corpus collossa, where subjects are causally influenced by phenomena that they report no conscious awareness of (Weiskrantz 1986), or in clinical breakdowns of the self, such as those involving "injected selves", or dissipated and disjoint mental lives (Stephens and Graham 2000). Clinical, medical, and cognitive psychology have represented distinct traditions studying mental pathologies, and as they begin to share more common phenomena, data, and theoretical bases, there is an opportunity for philosophers of psychology not only to contribute to discussions of foundational questions about the nature of the self, rationality, and normative mental functioning, but also to bring together these discussions with those on each of the three topics with which we began this review: mental representation, cognitive architecture, and consciousness.

Pain is the third and newest of these topics where philosophers and those in the cognitive sciences have begun to interact as a result of what we might call *phenomenal fallout*. The large community of researchers on pain have their home base in the medical sciences, and has focussed not so much on the theorization of pain as on its amelioration and treatment. Along with color, pain is the qualitative mental phenomenon most commonly invoked by philosophers discussing consciousness, and like color the empirical work on

pain has exploded in recent years. There are sensory and affective dimensions to pain, where the former reflects the role of pain as a detector of bodily damage, and the latter the phenomenal character of pain. Moreover, there turns out to be considerable interpersonal bodily variability for those experiencing pain. Conceptually, the sensory and affective dimensions to pain are distinct, and early empirical work offered support for the hypothesis that there are two separate pain systems. Dennett (1978) used some of the complexities to our common sense conception of pain to argue for an eliminativist view of pain, and more recently philosophers such as Valerie Hardcastle and Donald Gustafson have taken opposing views on whether pain is essentially perceptual or emotional in nature. My sense is that in the next few years the philosophical literature on pain will contribute significantly to ongoing debates generated by continuing empirical research (Aydede, Guzeldere, and Nakamura, forthcoming).

6. Embodied, Embedded, and Situated Cognition

A second general area in which there has been a hive of activity is that of *embedded* cognition, also referred to as *situated* or *embodied* cognition. In part as a reaction to the general character of traditional symbolic AI and connectionism, both of which abstracted away from the nature of the environment in which cognition actually operated, this cluster of views has emphasized the organism-environment coupling in theorizing about cognition. While the embedded movement is sometimes represented itself as anti-computational (e.g., Brooks 1997), there has been a concerted effort within an overarching computational framework to capture the spirit of the movement, ranging from Brian Cantwell Smith's (1996) reconceptualization of computation, to Daniel Dennett's (2000) emphasis on the important role of out-of-the-head scaffolding in higher mental processes. Central to the embedded movement is that idea that cognizers are *agents* who act in the world, gathering information about the world in order to so act. As this agent-centered conception of cognition has become increasingly a part of mainstream artificial intelligence (e.g., Russell and Norvig 1995), and indeed is one of the motivating themes of folk psychology, as I said at the outset, I view this development within the philosophy of psychology as less a departure from traditional views than a return to one of the themes familiar to those in the field (cf. Lakoff and Johnson 1999).

There is an obvious affinity between such approaches to cognition and the externalist views that have come to dominate philosophical reflection on intentionality and mental representation. There are a number of attempts (e.g., Clark 1997) to build some firmer bridges between the philosophical and scientific work, but in my view there is much more to be done here, in large part because I think that externalism about the mind has far more radical implications for the study of the mind than has been recognized (Wilson 2000, 2001).

see also: PHILOSOPHY OF BIOLOGY, PHILOSOPHY OF ECONOMICS, PHILOSOPHY OF PHYSICS, PHILOSOPHY OF THE SOCIAL SCIENCES

References

Aydede, Guzeldere, and Nakamura, forthcoming, *The Puzzle of Pain: Philosophical and Scientific Essays*. Cambridge, MA: MIT Press.

Block, N.,1995, "On a Confusion About a Function of Consciousness", *Behavioral and Brain Sciences*, 18:227-247.

Brooks, R., 1997, *Cambrian Intelligence*. Cambridge, MA: MIT Press.

Chalmers, D., 1996, *The Conscious Mind*. New York: Oxford University Press.

Clark, A., 1997, *Being There*. Cambridge, MA: MIT Press.

Crick F. and C. Koch, 1990, "Towards a neurobiological theory of consciousness". *Seminars in the Neurosciences* 2:263-275.

Cotterill, R., 1998, *Enchanted Looms*. New York: Cambridge University Press.

Dennett, D.C., 1978, "Why You Can't Make a Computer That Feels Pain", reprinted in his *Brainstorms*. Cambridge, MA: MIT Press

-----1987, *The Intentional Stance*. Cambridge, MA: MIT Press.

-----2000, "Making Tools for Thinking", in D. Sperber (ed.), *Metarepresentations*. New York: Oxford University Press.

Érdi P., 2000, "On the 'Dynamic Brain' Metaphor", *Brain and Mind*, 1:119-145.

Fodor, J.A., 1981, *Representations*. Sussex: Harvester Press.

-----1983, *The Modularity of Mind*. Cambridge, MA: MIT Press.

-----1987, *Psychosemantics*. Cambridge, MA: MIT Press.

Graham, G. and L. Stephens (eds.), 1994, *Philosophical Psychopathology*. Cambridge, MA: MIT Press.

Hardin, C.L., 1988, *Color for Philosophers*. Indianapolis: Hackett Publishing.

Haugeland, J., 1991, "Representational Genera", in W. Ramsey, S. Stich, and D. Rumelhart (eds), *Philosophy and Connectionist Theory*. Hillsdale, NJ: Erlbaum.

Horst, S., 1996, *Symbols, Computation, and Intentionality*. Berkeley, CA: University of California Press.

Jordan, M. and S. Russell, 1999, "Computational Intelligence", in R.A. Wilson and F.C. Keil (eds), *The MIT Encyclopedia of the Cognitive Sciences*. Cambridge, MA: MIT Press.

Kosslyn S., 1994, *Image and Brain*. Cambridge, MA: MIT Press.

Lakoff, G. and M. Johnson, 1999, *Philosophy in the Flesh : The Embodied Mind and its Challenge to Western Thought*. New York : Basic Books.

Levine, J, 1997, "Recent Work on Consciousness", *American Philosophical Quarterly* 34:379-404.

Lloyd, D., 2000, "Terra Cognita: From Functional Neuroimaging to the Map of the Mind", *Brain and Mind* 1:93-116.

Matthen, M., 1999, "The Disunity of Color," *Philosophical Review*,108:47-84.

Newell, A., and H. Simon, 1972, *Human Problem Solving*. Englewood Cliffs, NJ: Prentice Hall.

Pearl, J., 2000, *Causality: Models, Reasoning, and Inference.* New York: Cambridge University Press.

Port, R. and T. van Gelder (eds), 1995, *Mind as Motion*. Cambridge, MA: MIT Press.

Russell, S. and P. Norvig, 1995, *Artificial Intelligence: A Modern Approach*. Englewood Cliffs, NJ: Prentice Hall.

Searle, J., 1992, *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.

Smith, B.C., 1996, *On the Origin of Objects*. Cambridge, MA: MIT Press.

Stich, S., 1983, *From Folk Psychology to Cognitive Science*. Cambridge, MA: MIT Press.

Uttal, W., 2001, *The New Phrenology*. Cambridge, MA: MIT Press

Weiskrantz, L., 1986, *Blindsight*. Oxford: Oxford University Press.

Wilson, R.A., 1995, *Cartesian Psychology and Physical Minds*. New York: Cambridge University Press.

-----1999, "Philosophy", in R.A. Wilson and F.C. Keil (eds), *The MIT Encyclopedia of the Cognitive Sciences*. Cambridge, MA: MIT Press.

-----2000, "The Mind Beyond Itself", in D. Sperber (ed.), *Metarepresentations*. New York: Oxford University Press.

-----2001, "Two Views of Realization", *Philosophical Studies* 104:1-30.

-----2002, "Individualism", in T. Warfield and S. Stich (eds), *Blackwell Guide to the Mind*. New York: Blackwell.

Wilson, R.A. and F.C. Keil, (eds) 1999, *The MIT Encyclopedia of the Cognitive Sciences*. Cambridge, MA: MIT Press.