

ASKING WHAT'S INSIDE THE HEAD: NEUROPHILOSOPHY
MEETS THE EXTENDED MIND

Patricia S. Churchland, *Brain-Wise: Studies in Neurophilosophy*, Cambridge: MIT Press, 2002, xii+471, \$28.00, ISBN 0-262-53200-X (paper); \$70.00, ISBN 0-262-03301-1 (cloth).

John Bickle, *Philosophy and Neuroscience: A Ruthlessly Reductive Account*, Dordrecht: Kluwer Academic Press, 2003, xvi+235, \$37.00 1-4020-1302-7 (paper); \$96.00, ISBN 1-4020-7394-1 (hardback).

In their historical overview of cognitive science, Bechtel, Abrahamson and Graham (1999) describe the field as expanding in focus beginning in the mid-1980s. The field had spent the previous 25 years on internalist, high-level GOFAI (“good old fashioned artificial intelligence” [Haugeland 1985]), and was finally moving “outwards into the environment and downwards into the brain” (Bechtel et al, 1999, p.75). One important force behind the downward movement was Patricia Churchland’s *Neurophilosophy* (1986). This book began a movement bearing its name, one that truly came of age in 1999 when Kathleen Akins won a million-dollar fellowship to begin the McDonnell Project in Philosophy and the Neurosciences. The McDonnell Project put neurophilosophy at the forefront of philosophy of mind and cognitive science, yielding proliferating articles, conferences, special journal issues and books. In two major new books, neurophilosophers Patricia Churchland (2002) and John Bickle (2003) clearly feel this newfound prominence: Churchland mocks those who do not apply findings in neuroscience to philosophical problems as “no-brainers”; Bickle mocks anyone with traditional philosophical concerns, including “naturalistic philosophers of mind” and other neurophilosophers.

In this review, I will look at these books by Churchland and Bickle from the perspective of the *outward* movement of cognitive science: situated, embodied cognitive science. Such an analysis is especially important as proponents of the situated, embodied movement focus primarily on environmental information, perception, and action, and have (sometimes correctly) been taken to be no-brainers, denying the importance of the brain in understanding cognition. Given the ascendancy of neuroscience and neurophilosophy, such an attitude toward the brain simply will not do. The plan, then, is as follows. First, I will outline the different neurophilosophical programs, embodied in Churchland’s *Brain-wise: Studies in neurophilosophy* and Bickle’s *Philosophy and Neuroscience: A ruthlessly reductionist approach*. Then, I will recommend a way for situated, embodied cognitive scientists to come to terms with the undeniable importance of neuroscience and

neurophilosophy, showing that it is possible to resist the brain-obsessive wave without becoming no-brainers.

1. The first thing to say about these books is that they are both first-rate defenses of reductionism in the philosophy of mind. Churchland and Bickle each claim that the progress of the neurosciences shows, and seems poised to continue to show, that mental entities and processes are identical to brain entities and processes. They also agree that this fact has profound implications for philosophical questions. This is more or less the end of the similarity between the two books. I'll look at them in order.

Churchland's book is intended as a textbook, and it is an excellent one. It is introductory in that it presumes no background knowledge in philosophy or neuroscience. In some ways it is simply an update of *Neurophilosophy*, re-written to be easier going for students and to take account of the intervening advances in the neurosciences. It is also a widening of scope: where *Neurophilosophy* stuck more or less with the philosophy of science and philosophy of mind, *Brain-wise* uses neuroscientific evidence to address issues from all over philosophy. (The book is divided into three large sections, on metaphysics, epistemology, and religion.) Philosophically speaking, the new book maintains the pro-science, roughly pragmatist approach that made *Neurophilosophy* so popular with scientists. The main philosophical advance in *Brain-wise* is the inclusion of and reliance upon Rick Grush's emulation theory of representations (Grush 1997), more on which below.

As is appropriate in a textbook, Churchland is crystal clear about the nature of her approach, boiling it down to three hypotheses in the first chapter. The first hypothesis is the reductionist claim that mental activity *is* brain activity, and as such is to be studied by neuroscientists. This claim, of course, is the heart of Churchland's reductionism, and follows a clear, though highly simplified, account of what reduction is. The second hypothesis is that cognitive science, i.e., the science of mental activity, is important to neuroscience because it tells neuroscientists what the mental activity they need to account for while studying the brain. Cognitive science, that is, is the to-be-reduced theory. Churchland's third hypothesis is that to understand the mind one must understand the brain at all levels of organization. These hypotheses together give us a compelling recipe for conducting the traditional business of the philosopher. Start with a traditional philosophical problem. Be a good pragmatist/naturalized epistemologist and admit that the natural sciences are the best way to find a solution to it. If it is an epistemological problem, the answer is to be found in scientific studies of our minds as information processors, i.e., cognitive science. If it's a metaphysical problem, either it can be solved the same way or is pointless mystery-

mongering. Use results in cognitive neuroscience and/or brain-imaging to account for the results in cognitive science. Repeat, accounting for the cognitive neuroscience in terms of cellular neuroscience. And so on.

Churchland follows this recipe, more or less, throughout the book. Metaphysical fodder like the self, consciousness, and free will are accounted for in terms of information processing in the brain. Epistemology is naturalized, so epistemological problems come down to how the brain represents and learns. Religious experiences are cast as mis-interpreted brain squibs. My glib summaries don't do justice to Churchland's analyses of the issues, of course. Churchland has compiled an impressive pile of cognitive science and cognitive neuroscience to make these cases, and all of it is explained clearly and patiently. You will learn a lot of cognitive neuroscience reading this book, and you will not be bored learning it. Indeed, one good reason to pick up this book if you're a philosopher is for the cognitive neuroscience update. It is important, though, to notice the *more or less* in the first sentence of this paragraph. Churchland does follow the recipe for the first several steps, turning philosophical problems into scientifically-tractable problems concerning the brain's information processing. But, despite her hypothesis that it is important to understand the brain at all levels, she rarely touches on how the information processing she explains so carefully can be reduced to cellular activity or neurochemistry. These lower-level matters-the cellular, molecular and genetic-are Bickle's main focus.

Aiming to reduce the mental to the molecular is what makes Bickle's reductionism 'ruthless'. First, though, it is worth noting a more basic difference between Bickle's book and Churchland's. While Churchland's book is intended for everyone from advanced undergraduates on up, Bickle is writing for graduate students and professionals. Bickle's book is dense and, at times, difficult. This is, perhaps surprisingly, especially true when Bickle is writing about philosophy. Bickle's writing is patient and exceptionally clear when he explains the ins and outs of long-term potentiation at the cellular, chemical and genetic levels; he is less patient when discussing philosophical matters. Any reader, I think, could follow the neuroscience if they read carefully and, perhaps, have a neuroscience textbook handy. In contrast, only readers with graduate coursework in philosophy will follow all of Bickle's first chapter on the nature of reduction. There are two explanations for Bickle's differential presentation depending on the field of study. First, at the end of the day, the main audience for his book are professionals and advanced students in *philosophy*, along with a few philosophically sophisticated neuroscientists who are familiar with the details of the debate over reductionism. The second explanation is tied to Bickle's

main aim in the book, which is to argue that most work in philosophy of mind—including some of his own past work—really doesn't matter. Bickle even wonders at several points (e.g., p. 116) why the world needs philosophers of mind at all. According to Bickle's preferred approach to reduction, which he calls *new wave metascience*, we should simply take the word of neuroscientists when they claim to have reduced some bit of psychology to some bit of neuroscience. The change from being a philosopher of mind to being a new wave metascientist excuses one from worrying about replies to replies to replies to objections to theories of reduction. Indeed, one can stop worrying about any *a priori* barriers to reductionism, and trust neuroscientific reductionism-in-practice.

Looking at neuroscience as a new wave metascientist, Bickle finds that the reductionist action is not in cognitive neuroscience. Instead, he sees neuroscientists claiming to have reduced psychological processes directly to activity among genes and molecules, without an intervening reduction of the psychological to the cognitive neuroscientific. This is Bickle's ruthless reductionism: psychological facts, laws and generalizations are really molecular, neurochemical and genetic facts, laws and generalizations. Bickle exemplifies his "mind-to-molecules" reductionism most clearly with a detailed argument that the psychological phenomenon of memory consolidation (the production of long-term memories from short-term memories) *just is* long-term potentiation. Bickle's argument is based both on the consensus for such a reduction among neuroscientists and on an impressive recitation of the data. Bickle's wields this (purported) reduction with devastating effect against reductionist bugbears, especially for multiple realizability, widely held to be *the* reason that reductionism about psychology is implausible. Bickle begins his case against multiple realizability by changing the terms of the debate. As a new wave metascientist, Bickle is not interested in what psychology might be like on Mars or among angels; the new wave metascientist cares only about psychology in real creatures. Thus the debate shifts from whether the mind is multiply *realizable* in merely imagined creatures to whether it is multiply *realized* here on Earth. At this point, we simply do not know enough about either psychology or neuroscience to tell whether psychological phenomena in general are multiply realized, but we do know enough about some phenomena. Bickle shows in impressive detail that the very same molecules and mechanisms involved in long-term potentiation (and memory consolidation) in mammals are also at work in both houseflies and sea slugs. Bickle takes this truly surprising identity at the molecular level in such evolutionarily distant animals as evidence that long-term memory is realized in exactly one way on Earth. This is good evidence that long-term

memory, at least, is not multiply realized. The ruthless reductionist bets that the same is true of the rest of psychology.

This is exciting stuff, and in sharp contrast to the usual philosophical noodling and problematizing of the previously unproblematic. Churchland provides a plan for applying neuroscientific evidence to philosophical problems; Bickle applies it with a vengeance. Taken together, these books paint a picture in which real progress can be made on philosophical issues. Granted, the progress that Churchland and Bickle see is made by transforming the problems and then passing the buck to neuroscientists, and promising to put philosophers out of work. Despite the dimming job prospects for future philosophers of mind, it is hard to imagine preferring a static rumination over zombies, martian psychology and fading qualia to real answers. I, for one, am quite happy to join the battle against a priorism and endorse an increased role for empirical facts in philosophy. But, as I will argue in the next section, I think we need more facts.

2. As good and frequently convincing as these two books are in their advocacy of the downward expansion of cognitive science, the proponents of the *outward* expansion must, at least partly, resist their allure. For although most proponents of the situated, embodied cognitive science agree with Churchland and Bickle in their desire to apply results from the natural sciences to philosophical issues, they are not at all interested in reducing the mental to the neural. Indeed, this is the main point of situated, embodied cognitive science, which takes the object that requires explanation in psychology to be the combined brain-body-environment. What, then, should the situated, embodied cognitive scientist say about the undeniably good and important work done by Churchland, Bickle, and the neuroscientists whose experiments they cite?

The main strategy for the situated, embodied cognitive scientist here must be to insist the mental, though not magical or non-natural in any way, is not confined to the head. This leads to significant disagreements with reductionist neurophilosophers. First, for psychology to be reducible to neuroscience, the psychological is entirely in the head. Second, if the mental is all in the head, the way for the environment to impact the mind is by being mentally represented. The situated, embodied cognitive science proponent can deny both internalism and representationalism without deserving the epithets “no-brainer” or “a priori philosopher of mind”.

Since the 1970s, there have been a priori arguments that internalism about the mental is problematic. Putnam and Burge, for example, famously argued that “meaning ain’t in the head”. Given the scorn that Churchland and, especially, Bickle show for thought experiments,

we should not expect them to give up their internalism based on a priori considerations. Situated, embodied cognitive scientists, however, insist that psychology acknowledge a causal and explanatory spread outside the skull of the animal being studied. But to have a hope of convincing the neurophilosopher, or at least not provoking scorn, one must develop an *empirical* case against internalism. This has been the point of much work in situated, embodied cognitive science, whether by philosophers (Varela, Thompson and Rosch 1991, Clark 1997) or psychologists (Gibson 1979, Thelen and Smith 1994). Recently, Charles Heyser and I have provided evidence that some research in neuroscience itself suffers precisely because it is internalist. We found that researchers who use the object exploration experimental paradigm with rodents to study the effect of drugs, genes and neurotransmitters on behavior fail to take into account that the affordances of the to-be-explored objects, what actions the mice or rats can perform on them, dramatically affect the way rodents explore the objects. In many cases, experimentalists use objects with different affordances, e.g., some objects are climbable by rats, others are not. In such cases, the effects found in the experiments are just as likely to be due to differences among the affordances of the objects as they are to be due to the drug, gene or neurotransmitter under study. That is, it turns out that, at least in some cases, when neuroscientists, behavioral geneticists, and psychopharmacologists assume that studying the brain is studying the mind, they do flawed experiments or mis-interpret their results. (See Chemero and Heyser, to appear.) We argue that to fix these mis-interpretations, researchers must simultaneously investigate the brains, bodies and environments of their animals. This is also a problem for reductionism about the mental: if psychology is the study of the brain-body-environment, then neuroscience, which studies only what happens in the brain, is not a more general theory of the same subject matter as psychology. As such we should not expect a reduction of psychology to neuroscience.

As internalists, Churchland and Bickle are also representationalists. Bickle's representationalism strikes me as optional: if he thinks that the future of psychology and neuroscience is molecular, he has no need to worry about mental representations. And indeed, he spends very little time discussing them. Churchland's focus on cognitive neuroscience (i.e., to neural information processing), on the other hand, commits her strongly to representationalism. As such, issues of representation are at the forefront of her work, and explanations of how the brain might represent are crucial parts of her accounts of the self and of knowledge. As noted above, the main philosophical advance of *Brain-wise* over *Neurophilosophy* is in her adoption of Grush's emulation theory of representation. According to the emulation theory of representation

(Grush 1997), the most basic representations are forward models that provide the organism with a prediction of the outcome of its current action. Because they provide information about an animal's own body, emulators are especially appropriate for Churchland's discussion of self-knowledge, and she uses them to great effect there. Furthermore, there is evidence that something like emulators really are at work in the control of behavior in many species. (See Webb 2004, for example.) This might seem to be a problem for situated, embodied cognitive scientists, who are skeptical about value of internal representations and typically deny that representations are required in motor control. But it is far from obvious that emulators, though they may be ubiquitous in action control throughout nature, really are representations. Representations, traditionally, are about the external world, while emulators are control structures that are about the body. If one is not a very strong internalist, taking the mental to be confined to the central nervous system, one can quite naturally understand emulators as simply an element in the system they are controlling. Compare: the gears of a clock control the motion of its hands without being a representation of that motion. Emulators are similarly attached to the body parts they control.

Though I expect that Churchland, who clearly does feel that the mental is restricted to the central nervous system, would not accept this interpretation of emulators, she is open to the possibility that representationalism may not be the future of neuroscience. She mentions (though does not follow up) the possibility that dynamical systems theory might be a more appropriate language for neuroscience. Situated, embodied cognitive scientists ought to embrace this possibility. Dynamical systems theory is not only an appropriate language for understanding the activity of the brain, it is also the key to understanding the brain as part of a brain-body-environment system. It is quite natural for dynamical systems models to have some parameters that are in the brain, some that are in the body and some that are in the environment. Recent work by Bressler and Kelso (2001) and Varela and Thompson (2001) makes this point vividly. In both cases, dynamical systems models are shown to work both in brain-only explanations and in brain-body-environment ones.

3. Research in the neurosciences is progressing exceptionally rapidly. It is thus quite natural for philosophers like Churchland and Bickle to hope that neuroscientist can help us to provide solutions to many age-old philosophical problems. They express that hope enthusiastically in their very different new works. The strength of their convictions and their arguments can be intimidating to those of us who agree that brains are important, but don't think they are the whole story. I hope to have made it convincing that one can easily resist the pull of the

neurophilosophical wave without thereby being a “no-brainer”. To accomplish this feat, I offer the following three-step program for situated, embodied cognitive scientists: First, admit that brains are important; second, embrace dynamical systems modeling as the brain-friendly, but still non-internalist, means of explaining the activity of brain-body-environment; third, do not, under any circumstances, make arguments based on Twin Earth, inverted qualia, or martian psychology. The situated, embodied cognitive scientist can take results in neuroscience in stride by showing that they are best interpreted as shining light on a proper part of the larger brain-body-environment system.

ACKNOWLEDGEMENTS Thanks to Colin Klein reading a first-draft of this review. This paper was written while Tony Chemero was funded by National Science Foundation grant #00-04097

REFERENCES

- Bechtel, W., Abrahamsen, A., & Graham, G. (1998), The life of cognitive science. In W. Bechtel & G. Graham (Eds.), *A companion to cognitive science*. New York: Blackwell (pp. 1-104).
- Bickle, J. (2003), *Philosophy and Neuroscience: A ruthlessly reductionist approach*, Dordrecht: Kluwer Academic Press.
- Bressler, S.L. and Kelso, J.A.S. (2001), ‘Cortical Coordination Dynamics and Cognition’, *Trends in Cognitive Sciences* 5, -26–36.
- Chemero, A. and Heyser, C. (to appear), ‘Object Exploration and a Problem with Reductionism’, *Synthese*.
- Churchland, P.S. (2002), *Brain-wise: Studies in Neurophilosophy*, Cambridge: MIT Press.
- Clark, A. (1997). *Being There*, Cambridge: MIT Press.
- Gibson, J.J. (1979), *The Ecological Approach to Visual Perception*, Hillsdale, NJ: Lawrence Erlbaum.
- Grush, R. (1997), ‘The Architecture of Representation’, *Philosophical Psychology* 10, pp.-5–23.
- Thelen, E. and Smith, L. (1994), *A Dynamic Systems Approach to the Development of Cognition and Action*, Cambridge: MIT Press.
- Varela, F. and Thompson, E. (2001), ‘Radical Embodiment: Neural dynamics and Consciousness’, *Trends in Cognitive Sciences* 5, pp.-418–425.
- Varela, F., Thompson, E., and Rosch, E. (1991), *The Embodied Mind*, Cambridge: MIT Press.
- Webb, B. (2004), ‘Neural mechanisms for prediction: do insects have forward models?’ *Trends in Neurosciences* 27, pp.-278–282.

ANTHONY CHEMERO

Scientific and Philosophical Studies of Mind Program
Franklin and Marshall College
Whitely Hall
Lancaster, PA 17604, U.S.A.
and
Center for the Ecological Study of Perception and Action
University of Connecticut
406 Babbidge Rd, U 1020
Storrs, CT 06269, U.S.A.
E-mail: tony.chemero@fandm.edu

