Pergamon

Stud. Hist. Phil. Biol. & Biomed. Sci., Vol. 31, No. 4, pp. 615–627, 2000
Printed in Great Britain

www.elsevier.com/locate/shpsc

Essay Review

Contrary Views

*Nick Chater**

Paul M. Churchland and Patricia S. Churchland, *On the Contrary: Critical Essays, 1987–1997* (Cambridge, MA: MIT Press, 1998), xii + 349 pp. ISBN 0-262-03254-6 Hardback, £19.95.

Paul and Patricia Churchland are influential and controversial figures in the philosophy of mind and the theoretical foundations of the cognitive and neurosciences. *On the Contrary* brings together a selection of short and focussed pieces articulating and defending some of their most striking theses, many written as responses to critical attacks on their ideas. These pieces fit together to give a clear overview of the picture of the mind and brain that they advocate; all are sharply and entertainingly written, and tackle fundamental issues head-on. One of the theses in the book is the methodological precept, drawn from Feyerabend (1970), in favor of proliferating competing theories. The stimulating and radical positions that they have developed fully live up to this precept by widening the repertoire of contemporary discussions in philosophy, psychology and the neurosciences.

The Churchlands are self-proclaimed radicals, who intend to upset the *status quo* across a range of disciplines. Most of their critics, to whom their arguments in this book are addressed, are by comparison conservatives, arguing that the *status quo* is, in one way or another, more acceptable than the Churchlands would have us believe. While being a confessed radical has a pioneering and rather striking air, being a confessed conservative has less appeal. But, in reviewing this stimulating collection, this is the role in which I unwittingly find myself cast. Below, I will outline some of the Churchlands' core radical theses, and give some conservative rejoinders.

* Department of Psychology, University of Warwick, Coventry CV4 7AL, U.K.

PII: S1369-8486(99)00035-7

Before beginning, it is worth saying something about starting points. In weaving together threads from philosophy of mind and science with current developments in cognitive science and neuroscience, the question may arise: whose questions are being answered; and according to whose criteria? Nose-to-the-grindstone neuroscientists are likely to find much of the discussion of abstract theoretical and philosophical questions remote from their concerns; whereas philosophers may wonder just how relevant specific discussions of putative computational architectures of the brain are to traditional philosophical concerns. This is inevitable in interdisciplinary work of this kind; and such work provides a stimulus for those in different disciplines to perceive and develop links between disparate areas. Nonetheless, different readers will have very different biases and concerns in reading these essays, purely as a matter of disciplinary background. To declare my own starting point, so that my prejudices may be clear for readers of this review, my research area is cognitive science—developing and testing formal models of cognitive processes. As befits a cognitive scientist, and a conservative at that, I'll be concerned with philosophical issues primarily as they relate to general theoretical issues in developing an information processing theory of how the mind works; and I'll shy away from neuroscientific detail, aside from its significance for information processing analysis.

Let us turn, then, to consider some of the Churchlands' core theses. I'll divide the discussion into two groups of issues: those concerning folk theories and eliminativism; and those concerning the implications of neural network computation for theories of mental representation and the philosophies of mind and science.

1. Folk Theories and Eliminativism

Chapter 1 of *On the Contrary* opens with a beautifully clear outline of the historical origins of modern debates about the status of 'folk psychology', the rather nebulous and ill-defined cluster of competences, practices and beliefs which constitute our everyday mode of understanding each other's behavior. The Churchlands have argued that, construed as a scientific 'theory' of mind (as first bruited by Sellars, 1956), the folk theory, however it may be characterised, is a false theory, and should be rejected. Moreover, they argue, the notions which are central to the folk theory, such as belief, desire and the other propositional attitudes, are, thereby, on a par with terms in false theories in science. Thus *belief* should be treated in the same way as *phlogiston* or the *luminiferous ether*—there is *no such thing*. That is, folk theories are false, and the notions they postulated should be eliminated. Among other things, the Churchlands see folk theories as incompatible with current research on the computational properties of the brain, as we shall see below. The Churchlands endorse 'eliminative materialism'—the view that some materialistic theory of human minds, brains and behavior, presumably emerging out of the neurosciences, will ultimately replace folk psychology. This is a very radical conclusion to draw indeed: it requires signing up to the nonexistence of propositional

attitudes, (probably) the nonexistence of meanings for words and sentences, and the wholesale falsity of the underpinning explanation of human nature throughout the arts and the humanities.

A lot of attention has centered, in the critical attacks on the Churchlands to which much of this book is addressed, on whether it is really appropriate to say that folk psychology is a theory; and also on whether, if it is a theory, it is either appropriate, or even conceptually coherent, to declare that it is false. There are many deep issues here that I shall ignore in this review. But my hunch is that the Churchlands' view of folk psychology as analagous to, for example, alchemy is spot on. After all, both alchemy and folk psychology are loose collections of partially understood and mutually inconsistent beliefs and practices, which aim to systematize our understanding of aspects of the world, as well as, crucially, to control them. Folk psychology presumably functions to help gain control over the social world of other people, and perhaps also, indirectly, over oneself; alchemy aimed at the transmutation of substances, most critically the transmutation of base metals into gold. Moreover, in both folk psychology and alchemy knowledge is at best only loosely systematized and formalized. Indeed, neither domain has been formalized for, I would suggest, the same reason: the problems they attempt to solve (controlling the social world or transmuting lead into gold) were far too difficult to yield to theoretical analysis, given the conceptual tools available. Hence the attack on the problem had necessarily to be pragmatic and experimental, rather than formal and theoretical, if the enterprise was to get started at all. In both cases, too, the levels of prediction and control over the subject matter are consistently poor, although perhaps better than chance; and the theories are largely stagnant—they may mutate over time, but, by any outward measure of success, they progress little or not at all. I would therefore concur with the Churchlands that folk psychology is analogous to alchemy before chemistry, folk biology before modern biological science, astrology before astronomy, and so on.

Indeed, as the Churchlands argue, there seems nothing special about folk *psychology* in these regards; folk physics, medicine or meteorology seem to have the same characteristics. Considered as scientific theories, these parallel alchemy much better than they parallel chemistry. If we take it that alchemy is false, and chemistry is true, then this suggests that folk theories may have all the hallmarks of false theories. Similarly, if we take it that the ontology postulated by alchemy does not exist, whereas the ontology of chemistry does exist, then argument by analogy suggests that the ontologies of folk theories, across the board, do not.[1]

So let us accept, at least for the sake of exploring the conceptual territory, the falsity of folk psychology. But I suggest that the falsity of folk theories, including

[1]Of course, an antirealist might object that talk of theories being true, and the objects that the postulate existing or not, embodies a misleading picture of the issues. These deep matters need not concern us here, however—only whether folk theories are more analogous to systematic and successful science or to unsystematic and outmoded pre- or even pseudo-science.

folk psychology, has very different implications from those the Churchlands have advocated. I shall develop this viewpoint, and relate it to the Churchlands' point of view in three steps.

1.1. Folk Theories and Common Sense

We have talked so far about various aspects of common-sense knowledge that are loosely labelled 'folk psychology', 'folk physics', and so on. But common-sense knowledge of any kind has the same characteristics, whether that knowledge is of the natural world, the world of artefacts, or social and psychological matters. Consider, for example, the common-sense knowledge involved in knowing about household artefacts, such as chairs. Chairs are typically portable seats for one person; they are typically designed and constructed for this purpose; they are appropriate for a level rather than a sloping floor, and are shaped, to some degree, to fit the contours of the human body; they typically cost money and so on. These statements have rich interconnections with other aspects of common-sense knowledge. Thus, to understand what a seat is, what it means to be portable (given human lifting capacities), what it means to design an object, what the design significance of flat versus sloping floors is, what the rough contours of the (seated) human body are, what monetary transactions involve, and so on, requires drawing on a vast amount of further common-sense knowledge. Moreover, understanding each of *these* statements requires drawing on yet further common-sense knowledge, and so on indefinitely. The rich interconnectedness of common-sense knowledge is the primary reason that it seems inappropriate to conceive of common sense as organized into distinct and separate domains (see Putnam, 1988; Searle, 1992).

Such a view of common sense has important consequences for the Churchlands' arguments concerning the falsity of folk theories and the non-existence of the objects they postulate: first, that common-sense theories of *every* domain are false (at least on the assumption that they have exactly the characteristics that the Churchlands identify as damning for folk psychology); and second, and more radically, that the terms of common-sense ontologies quite generally refer to things that do not exist. In particular, categories such as *chair*, *tennis*, and *jazz*, or indeed those associated with any non-technical content word in the dictionary, do not exist—these terms do not refer to anything! Exactly the same arguments that show the status of *belief* to be analogous to the status of phlogiston rather than that of oxygen do exactly the same job for any common-sense term.

The Churchlands are thus caught in a dilemma. On the one hand, they might simply accept the above conclusion, despite its startling and counterintuitive character. For example, they might argue that our intuition that chairs exist is really an intuition concerning particulars—that there are particular things (such as those tucked under my dining table), which are commonly labelled chairs, and that these particulars exist. But the radical claim is not about particulars at all, but universals—it is about the viability of the category of chairhood. To develop this point

with an analogy, suppose that we were to announce in medieval Europe that there are no witches. The response would be: 'But what about the women we drowned last month, with the black cats and the broomsticks? Are you telling us they were figments of our imagination?' But the claim about witches is also a claim about universals—that there is no viable category of witch-hood. Of course, the particular people commonly labelled as witches do exist. Equally, perhaps, our intuitive outrage over the claimed non-existence of chairs should be given little credence. Many philosophers and scientists will not be reassured—a world in which common-sense theories are false and common-sense categories are incoherent, across the board, may seem unacceptable. As it happens, I think this way out of the dilemma is the right one (see Chater and Oaksford, 1996, for a defense of this view).

The alternative way out of the dilemma is to argue that folk psychology, folk physics and folk medicine are fundamentally different from other aspects of common-sense knowledge. In Chapter 2 of *On the Contrary* Paul Churchland advocates this line with respect to so-called 'functional kinds', such as *chair*. He argues that the difference between legitimate functional kinds and illegitimate folk psychological categories stems from the fact that 'the physical tokens of any functional kind are typically manufactured to meet our functional specifications and typically there is no intelligible question of whether our functional concept is adequate to the behavioral reality the manufactured object displays' (p. 27). So functional kinds are grounded in the specifications of the manufacturer; whereas terms in folk psychology and other folk theories are grounded, if they are grounded at all, by correctly describing the natural world.

But this defense seems inadequate. The problem is that the divide between artefacts and natural phenomena is actually not as clearcut as the Churchlands' picture suggests. For example, consider linguistic, economic or collective social phenomena which we might describe with folk categories such as *word*, *sentence*, *money*, *loan*, *democracy*, or *law*. These are human constructs—and thus in a sense they are artefacts of human culture; but their creation does not follow the pattern in the Churchlands' sketch. No-one designed such phenomena; and hence there is no question that such constructs gain their coherence simply in virtue of fitting, by fiat, the designer's specification. Moreover, categories such as those we have listed seem to have deep theoretical roots—folk explanations of linguistic, economic or sociopolitical phenomena seem to be as much embedded in our notions of how language, finance or society works as are folk explanations of psychological, biological, or physical phenomena. Thus, just because an aspect of the world is of human creation does not imply that it is designed—and understanding such aspects of the world appears to be a theoretical enterprise, just as biology, chemistry or psychology are theoretical enterprises. But, on reflection, this point applies even to apparently prototypical examples of 'designed' products. Consider a category such as *chair*, which the Churchlands would like to save as a viable functional category. This category is not specified by a chair designer, who cannot legislate

for what counts as a chair or not. Instead, the emergence of the production and use of artefacts such as chairs is merely another complex emergent social phenomenon, outside of any individual designer's control. This is all too evident in legal contexts, where attempting to specify what counts as a 'chair,' a 'car' or a 'road' involves an indefinitely large amount of open-ended analysis (see Hahn and Chater, 1998, which attempts to explore the interconnections between artefact categories are related notions (e.g., roads are thoroughfares along which motor vehicles can pass; they are typically of a certain width; their surfaces have certain properties; and so on)). In short, understanding artefacts, like other cultural products, is a matter of theory in just the same way that understanding the natural world is; and hence if terms in folk theories of the natural world are incoherent ('beliefs do not exist'), then the same applies to terms in folk theories of artificial domains ('money does not exist', 'chairs do not exist', and so on).

In sum, I believe that the Churchlands do a good job of persuading us that common-sense theories are more analogous to disreputable, rather than reputable scientific theories. But this conclusion applies across the board, to common-sense knowledge of any aspect of the world; and hence undermines not just categories such as *belief* but everyday categories quite generally.

1.2. Folk Psychology and Cognitive Science Concern Different Domains

The Churchlands see folk theories as displaced by scientific theories; and they see folk psychology as liable to be displaced by a scientific account of mind. Taking a broad construal of folk psychology, to include commonsensical ideas of any and every kind concerning how the mind works, this is reasonable enough. For example, everyday views about how vision works have clearly been overtaken by the developments in visual science; and everyday ideas about the structure and processing of language have been superseded by linguistics, psycholinguistics and computational linguistics.

But what philosophers, including the Churchlands, typically focus on as the core of folk psychology is more constrained. This is the body of knowledge which underlies everyday explanations of human behavior in terms of beliefs, desires and other propositional attitudes. Such explanations are both extremely widespread and of enormous importance. Propositional attitudes allow us to ascribe meaningful mental states to each other, to explain thought in terms of internal transitions between meaningful states, and to explain how mental life is embedded in the world, by connecting the contents of these internal states with perception, action and language.

By saying that a person has a particular propositional attitude, we are ascribing to that person a mental state which has the same content as the 'proposition' component of a sentence of natural language expressing that propositional attitude. Thus, propositional attitude explanation recruits all the resources of natural language and allows them to be used to characterize mental states and their transitions.

The central idea is that the internal mental state transitions correspond to natural language *arguments*, with the relationship between mental states and natural language being given by the ascription of propositional attitudes. Rational thought corresponds to good argument, and irrational thought to flawed or incoherent argument. People are rational to the degree that their behavior can be reconstructed in terms of the outcome of rational arguments.

The Churchlands suggest that this style of explanation will gradually be overtaken by work in the cognitive and neurosciences—indeed, they suggest that the process of overthrow may already be underway.

I believe that this appearance is mistaken. Pickering and Chater (1995; see Morris and Richardson, 1995, and Chater and Pickering, 1997, for replies) argue that the cognitive and neurosciences have not even begun to encroach on the territory of folk psychology, narrowly construed. *De facto*, the materialist scientific study of mind and folk psychology deal with completely non-overlapping aspects of mental life. Specifically, folk psychology applies to 'knowledge-rich' aspects of cognition, which have proved completely unamenable to the cognitive and neurosciences. In artificial intelligence, researchers refer to avoiding the world-knowledge problem—if a proposed solution to a computational problem involves drawing on everyday common-sense knowledge, then it can immediately be rejected as currently unsolvable. Moreover, as we have already seen, knowledge-rich processes cannot be understood piecemeal: attempting to formalise just a part of world knowledge leads to having to spell out all the background knowledge that it presupposes, and so into the indefinite and currently entirely unfeasible program of capturing world knowledge in its entirety (this is what Fodor, 1983, refers to as the *isotropy* of common-sense knowledge). Because knowledge-rich processes have been so difficult to analyse, progress in the cognitive and neurosciences has been limited to 'knowledge-free' aspects of cognition, about which folk psychology is silent.

So, for example, cognitive science is successful at understanding relatively low-level perception and language processing, motor control, the structure of memory stores, the limits of attention, the structure of concepts, and so on. Theories in these areas make no reference to beliefs or desires, because they fall into one of two categories. Either they are concerned with special-purpose, modular, psychological processes, which are isolated from world knowledge; or they are concerned with structural aspects of the mind, rather than the content of what is represented. By contrast, the domain of folk psychology is the contents of everyday, common-sense thought, which is paradigmatically knowledge-rich. But cognitive science and the neurosciences have been conspicuously unsuccessful in unravelling this aspect of mental life. For example, the present hopelessness of formalizing common-sense knowledge is evident in the notorious 'frame problem' in artificial intelligence (see McCarthy and Hayes, 1969; Pylyshyn, 1987).

This viewpoint suggests that folk psychology and cognitive science should not be seen as standing in competition, with folk psychology liable to succumb to the

march of scientific progress. Instead, the two programs have entirely different subject matters, because the cognitive and neurosciences simply cannot handle the knowledge-rich phenomena which are the subject matter of folk psychology. On the present course, it seems unlikely that folk psychology, in the narrow sense of explanation which relies on the semantic content of propositional attitudes, will be eliminated or even encroached on, by the cognitive and neurosciences—at least, as long as the world-knowledge problem remains unsolved.

1.3. Why Eliminative Materialism?

A final issue concerns why the Churchlands insist that folk psychology, along with other folk theories, should be eliminated in terms of *materialist* theories. It is hard to tell from this book whether the 'materialism' that the Churchlands advocate is entirely uncontroversial, or whether it is a strong and radical thesis. They reject type–type reduction between explanatory levels—therefore they do not expect that economic phenomena, for example, can be re-expressed in the vocabulary of the physical sciences. But if 'materialism' in this context comes to no more than a commitment to token–token reduction (and, in the context of philosophy of mind, to a rejection of dualism), then the thesis is uncontroversial and uninteresting.

In particular, materialism, in this weak sense, is quite compatible with the possibility that folk theories might be replaced with better 'scientific' theories *at the same level of abstraction*. Indeed, this has occurred in chemistry and biology—notions such as phlogiston have been replaced by notions such as oxygen, and folk classifications of the biological world have been replaced with scientific taxonomies. Crucially, the replacement notions seem no more materialist (in any sense the term may have) than the folk notions that they replace. Similarly, one might expect that a scientific substitute for folk psychology might be pitched at roughly the same level of analysis; or at least it would seem that this is entirely plausible. Unless this possibility is ruled out, there would seem to be no substantial implications of the falsity of folk psychology with respect to questions of materialism. Instead, however, the Churchlands are committed to the expectation that the replacement will be unremittingly neuroscientific—it will concern the mechanistic structure of the brain. This is a radical and crucial step, but it appears to slip by unnoticed and unargued.

2. Philosophical Implications of Neural Network Computation

I have so far ignored the specific discussions of neuroscience and neural network computation that run through *On the Contrary*, and I have considered philosophical issues in isolation. But this gives only a partial picture of the Churchlands' philosophical strategy, because considerations of progress in the cognitive and neurosciences are central to their philosophical arguments.

Here, the conservative cognitive scientist in me will be especially to the fore. I believe that the picture of progress painted by the Churchlands is excessively rosy.

For example, in Chapter 4 Paul Churchland cites the literature on neural networks as providing an alternative to the classical propositional conception of mind. But, in practice, neural networks have addressed problems traditionally viewed as requiring propositional representations by providing novel implementations of standard symbolic propositional representations (Derthick, 1987; Shastri and Ajjanagadde, 1993; Smolensky, 1990; Touretzky and Hinton, 1988). Otherwise, neural network researchers have steered clear of such issues, and instead concentrated on developing highly specific cognitive or biological models in domains which are not touched on by traditional propositional attitude explanation. By contrast, the Churchlands talk as if particular neural networks provide a theory of how the mind works (see, e.g., Chapter 11 on recurrent networks as models of cognition). The computational capacities of recurrent neural networks, at least as currently studied, are, however, severely limited—such a network might be a plausible, partial model of, for example, single word reading (see Plaut *et al*., 1996), or even a radically simplified account of some aspects of parsing (see Christiansen and Chater, in press); but they are far too limited to serve as new models of the mind as a whole (especially given the difficulty of the world-knowledge problem discussed above).

Other attempts to draw philosophical implications from neural network computation seem puzzling. For example, in Chapter 15, Paul Churchland argues that neural network models are significant in relation to some of Feyerabend's theses in the philosophy of science—this is part of a more general reconstruction of epistemology from a 'neurocomputational' perspective. He argues, for example, that in a neural network model, the theory-ladenness of observation arises automatically. '. . . [A]ny configuration of synaptic weights dictates a specific set of partitions on the activation space of the postsensory neurons to which they connect. And that set of partitions constitutes a specific conceptual framework or theory . . . no cognitive activity takes place save as the input vectors pass through that speculative configuration of synaptic connections, that theory. Theory-ladenness . . . makes processing activity genuinely *cognitive* in the first place' (pp. 270–271). But neural networks can equally well be constructed to provide a theory-observation dichotomy. For example, many neural network models assume that perceptual processing is strictly distinct from general knowledge, including 'theories' in the normal scientific sense—for example, perceptual processing may be arranged in 'modules'. In this framework, the output of the perceptual system might be viewed as corresponding to theory-independent observation—this view is, for example, pursued by Fodor (1983). More generally, many theorists argue that perception, whether strictly modular or not, is cognitively impenetrable, that is, cannot be influenced by general knowledge, including knowledge about scientific theories (see Pylyshyn, 1984, in press). The conclusion is, then, that the proposal that cognition should be understood in terms of neural networks, rather than in terms of symbol processing, seems to be independent of the question of whether perceptual observation is or is not theory-laden.

Nonetheless, the general emphasis here is that studying current computational models of learning, especially those potentially relevant to what may occur in the brain, is important for epistemology and the philosophy of science. Churchland stresses that neural networks provide a very different perspective on what it means to acquire new knowledge from the perspective implicit in a classical view, where propositional attitudes are central. On the latter view, observations and hypotheses are postulated in a propositional language, and learning involves inferential operations, in line, at least approximately, with some normative theory of inference (such as logic, probability and decision theory). But in neural networks, knowledge is represented in patterns of connections in a network, and learning involves updating those connections, typically in response to simple local rules concerning the activity of the 'nodes' or 'neurons' which make up the network.

These pictures are interestingly different, but before jumping to the conclusion that neural networks show how to overthrow the classical epistemological picture, a few caveats are in order.

First, as we have noted, neural networks are extremely limited computationally. The Churchlands may feel that these limitations are illusory, first because 'neural networks have been shown to be "universal approximators" (Hornik *et al*., 1989) . . . [and] such networks can also *learn* to approximate any desired function, from repeated presentation of its instances, by means of various automatic learning procedures' (p. 41, emphasis in the original). But the picture is not as reassuring as this might suggest. The first result is analogous to saying that Fourier series can approximate arbitrary functions.[2] The second result is analogous to saying that Fourier coefficients can be estimated from empirical data. Although true, these results are not reassuring in practice, because it is unclear to what class of problem neural networks, or Fourier series, or any other universal mathematical approximation method, can actually be usefully applied in practice. In reality, it is all too easy to choose functions (e.g., mapping from a speech wave to a representation of the meaning of what is said; mapping from images to descriptions of the scenes they depict; and so on) which prove entirely intractable to general-purpose neural network methods. All the considerations of cognitive architecture and complex structured representation which elementary neural networks do away with appear to be necessary even to begin to tackle difficult cognitive problems of this kind. To caricature the situation, a parallel to the Churchlands' point of view would be—'we don't need all these high-level programming languages, data-bases, or graphics packages, or complex data structures and algorithms of any kind; after all, all digital computation can be done with networks of binary transistors; that's all we need to know about'. Of course, at some level a digital computer is just a network of transistors; and a brain is just a network of neurons. But understanding both

[2]Both neural network and Fourier results apply only if the function obeys certain technical conditions, of course.

computational systems may also require description at higher levels, which may very well require postulating structured propositional representations.

The second caveat regarding the epistemological implications drawn from neural networks as models of thought is that the link between neural networks and the classical propositional picture is perhaps tighter than might be imagined. Most successful neural network learning systems can be interpreted in probabilistic and statistical terms—typically network computation can be viewed as instantiating some, perhaps rather complex, probabilistic model; and learning involves statistical inference concerning the adjustment of the parameters of that model (Chater, 1995; McClelland, 1998; MacKay, 1992). This relationship is not merely a mathematical curiosity: it helps explain when and why neural networks learn successfully; it has also been crucial in guiding the development of novel learning algorithms. Moreover, there have been some important links found between an important class of symbolic probabilistic reasoning systems, belief networks (Pearl, 1988) and a class of neural network models, based on the asymmetric Boltzmann machine (Neal, 1992).

The final caveat is that, despite the Churchlands talking freely of 'state–space semantics' in relation to the meaning of vectors of network activation, there is currently not even the ghost of a theory of meaning for activation vectors or synaptic connection strengths in a neural network. The only theories of meaning we currently have available are defined over propositional representations—it is not even clear that it makes sense to conceive of non-propositional structures as representations at all. There are deep and unsolved philosophical issues here, of course, concerning the very nature of meaning, truth and representation. But I want to stress instead the computational issue: from a computer science perspective we simply have no idea how to assign meanings to activations patterns (e.g., of hidden units in a back-propagation network), or still less to the patterns of connectivity in a network (which, in Chapter 15, are identified with an agent's 'theory' of the world). Given this state of affairs, the Churchlands' upbeat talk of a new 'semantics for cognitive activity' (p. 41) as a present reality is off the mark.

3. Conclusion

On the Contrary is a stimulating and entertaining read. Its authors are, very laudably, not interested in searching for a middle road—they have taken a series of radical theses, and pushed them to their logical, and perhaps disconcerting conclusions. Few will agree with everything they argue for; and many will disagree with a great deal. Just as most of its chapters arose from heated debate, so this volume itself will generate further discussion and controversy. I suspect this is exactly what its authors had in mind.

References

Chater, N. (1995) 'Neural Networks: The New Statistical Models of Mind', in J. P. Levy, D. Bairaktaris, J. Bullinaria, and P. Cairns (eds), *Connectionist Models of Memory and Language* (London: UCL Press), pp. 207–227.

Chater, N. and Oaksford, M. (1996) 'The Falsity of Folk Theories: Implications for Psychology and Philosophy', in W. O'Donahue and R. F. Kitchener (eds), *The Philosophy of Psychology* (London: Sage), pp. 244–256.

Chater, N. and Pickering, M. (1997) 'Two Projects for Understanding the Mind: A Response to Morris and Richardson', *Minds and Machines* **7,** 553–569.

Christiansen, M. and Chater, N. (in press) 'Toward a Connectionist Model of Recursion in Human Linguistic Performance'. *Cognitive Science* **23**, 157–205.

Derthick, M. (1987) 'A Connectionist Architecture for Representing and Reasoning about Structured Knowledge', Research Report No. CMU-BOLTZ-29, Department of Computer Science, Carnegie Mellon University, Pittsburgh.

Feyerabend, P. (1970) 'Against Method: Outline of an Anarchistic Theory of Knowledge', in M. Radner and S. Winokur (eds) *Analyses of Theories and Methods of Physics and Psychology*, *Minnesota Studies in the Philosophy of Science* (Minneapolis: University of Minnesota Press), pp. 1–130.

Fodor, J. A. (1983) *Modularity of Mind*. (Cambridge, MA: MIT Press).

Hahn, U. and Chater, N. (1998) 'Similarity and Rules: Distinct? Exhaustive? Empirically distinguishable?', *Cognition* **65,** 197–230.

Hornik, K., Stinchcombe, M., and White, H. (1989) 'Multilayer Feedforward Networks Are Universal Approximators', *Neural Networks* **2,** 359–366.

MacKay, D. J. C. (1992) 'A Practical Bayesian Framework for Backpropagation Networks', *Neural Computation* **4,** 448–472.

McCarthy, J. M., and Hayes, P. (1969) 'Some Philosophical Problems from the Standpoint of Artificial Intelligence', in B. Meltzer and D. Michie (eds) *Machine Intelligence* (Edinburgh: Edinburgh University Press), pp. 463–502.

McClelland, J. L. (1998) 'Connectionist Models and Bayesian Inference', in M. Oaksford and N. Chater (eds), *Rational Models of Cognition* (Oxford: Oxford University Press), pp. 21–53.

Morris, A. and Richardson, A. (1995) 'How *Not* to Demarcate Cognitive Science and Folk Psychology: A Response to Pickering and Chater', *Minds and Machines* **5,** 339–355.

Neal, R. M. (1992) 'Asymmetric Parallel Boltzmann Machines Are Belief Networks', *Neural Computation* **4,** 832–834.

Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (San Mateo, CA: Morgan Kaufman).

Pickering, M. and Chater, N. (1995) 'Why Cognitive Science Is Not Formalized Folk Psychology', *Minds and Machines* **5,** 309–337.

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., and Patterson, K. (1996) 'Understanding Normal and Impaired Word Reading: Computational Principles in Quasi-Regular Domains', *Psychological Review* **103,** 56–115.

Putnam, H. (1988) *Representation and Reality* (Cambridge, MA: MIT Press).

Pylyshyn, Z. W. (1984) *Computation and Cognition: Toward a Foundation for Cognitive Science* (Montgomery, VT: Bradford).

Pylyshyn, Z. W. (ed.) (1987) *The Robot's Dilemma: The Frame Problem in Artificial Intelligence* (Norwood, NJ: Ablex).

Pylyshyn, Z. W. (in press) 'Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception', *Behavioral and Brain Sciences*.

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

Sellars, W. (1956) 'Empiricism and the Philosophy of Mind', in H. Fiegl and M. Scriven (eds) *Minnesota Studies in the Philosophy of Science* (Minneapolis: University of Minnesota Press).

Shastri, L. and Ajjanagadde, V. (1993) 'From Simple Associations to Systematic Reasoning:

A Connectionist Representation of Rules, Variables, and Dynamics Bindings Using Temporal Synchrony', *Behavioral and Brain Sciences* **16,** 417–494.
Smolensky, P. (1990) 'Tensor Product Variable Binding and the Representation of Symbolic Structures in Connectionist Systems', *Artificial Intelligence* **46,** 159–216.
Touretzky, D. D. and Hinton, G. E. (1988) 'A Distributed Connectionist Production System', *Cognitive Science* **12,** 423–466.