

IS 'MIND' A SCIENTIFIC KIND?¹

Andy Clark

1. Three models of mind and cognition

The title question (Is 'Mind' a scientific kind?) invites a consideration of just about every major problem in Philosophy of Science and several in Philosophy of Mind. Needless to say, I do not propose to attempt anything quite so grand. Instead, I will really address the somewhat narrower question: how should we conceive the relation between scientific studies of cognition and the folk ontology which depicts minds as loci of beliefs, desires, concepts, propositional attitudes etc.? In particular, I shall first consider and reject two extreme options, viz:

(a) That the folk ontology must, on pain of Eliminativism, be reconstructible using only the resources of some scientific study of cognitive processes.

(b) That the folk ontology is legitimated by gross behaviour patterns alone and is conceptually independent of whatever science can tell us about inner states and processes.

As a kind of rough shorthand, I shall describe position (a) as the thesis that the various items in the folk ontology, to be real, must turn out to name inner scientific kinds. And I shall describe (b) as the thesis that such items name purely observational kinds. The rejection of (a) will flow from a discussion of the point and purpose of folk psychological talk. The rejection of (b) will flow from a discussion of familiar counter-examples to the thesis that mind is just an observational kind viz. another look at the contemporary Cartesian demons: Giant Look-up Tables and Quantum Fluke Beings. I shall end by developing an alternative thesis:

(c) That the folk ontology of minds and mental contents, although not required itself to pick out scientific kinds, is nonetheless required to be

broadly *intelligible* given a correct scientific understanding of cognition.

The challenge, of course, is to make clear what this notion of intelligibility amounts to insofar as it is something weaker than the requirement that folk constructs must name scientific kinds. I take some steps in that direction by introducing requirements concerning the intelligibility of depicting a system as issuing recall-dependent judgements and the intelligibility of depicting it as a locus of conscious mental states and/or qualitative experience. I suggest that our intuitions about the contemporary demons (Giant Look-up Tables etc.) are best treated as rooted in such requirements.

In sum, I shall argue that the folk ontology of mind is in essence a practical tool which makes minimal (but real) demands on the types of inner workings compatible with its correct deployment. The product is thus a version of Dennett (1987), but one which tries also for a concrete picture of some minimal inner requirements on True Believers. Or, if you prefer, it is a version of Ryle (1949), but without the total rejection of science as impacting upon the conception of mind.

2. *Against Super-Fodorian realism*

The thesis that the folk ontology of mental states must, on pain of Eliminativism be neatly reconstructible in some more scientific milieu amounts to a doctrine which I label Super-Fodorian Realism (Clark, 1993). The doctrine is *super-Fodorian* in that where Fodor sees the existence of a folk-content encoding inner code as an empirical fact (the a posteriori explanation of the systematicity of thoughts ascribed using the apparatus of folk psychology), the Super-Fodorian sees the existence of such a code as conceptually essential to the truth of belief/desire citing explanations. If science were to show us that *concepts* and *propositions* do not exist as scientifically identifiable inner items the folk mentalistic ontology would (according to the Super-Fodorian) be bankrupt. Why should this be so? Although there exist a variety of Super-Fodorian arguments in the contemporary literature, they all share a basic structure. I shall first exhibit that structure, then flesh it out with a single representative example. The basic structure involves an argument from *disunity*. It goes like this:

2.1 Argument from disunity

1. Folk psychology individuates mental states using an apparatus of attitudes (belief, desire, etc.) and propositions. And it individuates propositions as distinct structures of concepts.
2. Folk psychology is thus committed to a specific account of *sameness for mental states*, viz. that distinct mental states may involve different attitudes to the *same* propositional content, and that different propositional contents may involve the *same* concept.
3. But suppose a good scientific story about our inner cognitive workings fails to identify scientifically respectable inner states which recapitulate these judgements of sameness? Suppose the inner story posits unstable, elusive or fragmentary items where the folk story posits a recurrent entity (concept or proposition)?
4. In such cases the folk story must be abandoned as the order it depicts is revealed as illusory.

Versions of the argument from disunity can be found in Davies (1991), Stich (1983) and Ramsey, Stich and Garon (1991). To report a single example, Davies (1991) depicts the folk as committed to a vision in which an individual's mastery of a given concept (say, the concept of 'bachelor') is invoked to explain a host of behaviours (e.g. all their inferences from the information that so and so is a bachelor to the conclusion that so and so is unmarried). The folk thus invoke a single item (the concept 'bachelor') in a variety of explanations of someone's behaviour (verbal behaviour, in this case). But, Davies insists, the mere fact that there is a discernible pattern in the individual's inferential behaviour cannot guarantee that it is mastery of a single concept which explains the behaviour. What if our hero/heroine is a big look-up table with a separate entry for the inference for each possible name in the language? The folk story, in discerning an underlying unity in the observed behaviours, is committed (Davies argues) to a certain kind of inner scientific story viz. one in which the concepts picked out by folk-psychology exist as discrete and literally recurring inner syntactic items. In short, there had better exist something very like an inner language of thought if the folk explanations are to be accepted as legitimate.

The trouble with such Super-Fodorian arguments is that they trade on ambiguities in the notion of sameness. In Davies' case, the ambiguity is between the claims

(a) that the same *concept-mastery* is implicated in several behaviours and

(b) that the same concept (conceived as a kind of discrete inner data-structure) is present and active in several behaviours.

Thus suppose that learning to use the word 'dog' in a way which meets public criteria involves training several disparate and internally disunified cognitive resources. What the training results in is thus a kind of tuning of many different parts of an overall system. Upon successful completion of such training, we say of someone that he has mastered the concept. If on one occasion he then uses inner resource X to power an appropriate response and on another occasion he uses a different inner resource, Y, it remains true to say that it is, in a sense, the single *concept-mastery* which explains each behaviour. Yet it is also true that (in Davies terminology) there need be no causal common factor active on each occasion. Grasp of a concept, I therefore want to say, may be akin to possession of a *global skill* (cf. Evans (1982); pp. 101-102). Just as an individual may be said to have a skill at golf (a skill which explains both successful putting and successful driving) and yet deploy quite distinct cognitive sub-skills to power various manifestations of this global skill, so she may possess global conceptual skills whose internal cognitive underpinning is various and fragmentary.

Such an image (of folk psychology as naming global skills emergent out of potentially messy and disunified complexes of inner workings) allows folk-individuative practices to co-exist with several recent lines of scientific conjecture concerning cognition, viz:

1. The evidence of internal disunity between brain systems responsible for verbal and non-verbal behaviours (used misguidedly by Stich (1983) as an argument in favour of Eliminativism)
2. The evidence for an unexpectedly strange and rich body of possible dissociations of cognitive abilities coming from cognitive neuropsychology (Ellis and Young, 1988; Warrington and McCarthy, 1990; Shallice 1988; Humphreys and Riddoch, 1987)
3. The distributed connectionist model of lexical knowledge in which local contextual information results in subtly different internal representations corresponding to the same folkindividuated content on different occasions (cf. Elman (1991)'s comment that in his lexical categorization network there are no recurrent canonical representations of lexical items and that instead "it is literally the case that every occurrence of a lexical

item has a separate internal representation").

In the inner economy, the folk items may well dissolve, reform or fragment. Things fall apart. But if our description of the folk vocabulary as specifying global skills evidenced in the daily behaviours of whole agents is correct, then such fragmentation, in and of itself, is interesting but harmless. The folk mentalistic ontology need not specify integral scientific kinds. It is an easy mistake, however, to think that the folk ontology, if it is not required to name scientific kinds, is altogether immune to the influence of scientific discoveries. Not so, as we shall now see.

3. *Against mere ascriptivism*

In attempting to insulate folk psychology from the misguided demands of the Super-Fodorian realist, it is tempting to fall back on a position in which all that matters, as far as the acceptability of the folk ontology and explanations goes, are the patterns in gross behaviour. The presence of these patterns is then seen as exhausting the commitments of the folk discourse. Where such patterns exist, according to this vision, the folk discourse is properly deployed. The most famous (albeit most slippery) proponent of such a pure ascriptivism is the Daniel Dennett of Dennett (1987,1988). Dennett calls the practice of explaining something's behaviour by reference to beliefs/ desires etc. "taking the intentional stance" towards that object. And he comments that: Any object--whatever its innards--that is reliably and voluminously predictable from the stance is in the fullest sense of the word a believer (Dennett, 1988,p.496):

"Any object -whatever its innards- that is reliably and voluminously predictable from the stance is in the fullest sense of the word a believer"

The trouble with such a forthright ascriptivism, famously, is that it threatens to let in bizarre cases in which we have strong intuitions that, voluminously predictable from the intentional stance or no, the object in question just ain't a True Believer (a haver of genuine beliefs and desires). A classic case is the Giant Look-up Table: a super-fast, super-large computer which stores a distinct output for each and every one of a vast

number of inputs and uses this brute force approach to produce behaviour in which Dennett style patterns are rife. Or you may prefer my own example, the Quantum Fluke Being: a cosmic accident creature which gets all the behaviours right but does so by an increasingly unlikely (but never 100 per cent impossible) series of accidents. Its innards are disorganized mush, yet it exhibits nice patterns in gross behaviour. Or you may prefer Lycan's Zombies (Lycan, 1988, pp.518-519), or Chris Peacocke's Martian Marionettes (Peacocke, 1983). The moral is the same: the folks' commitments don't stop at the surface of the skin.

So where do they stop? Davies had a neat story (Section 1 above) in which the lack of re-usable syntactic entities corresponding to the semantic items of the folk story torpedoed look-up tables and their ilk. But we found it too demanding. Bennett (1991) offers, interestingly, an almost diametrically opposed thought, viz. that genuine intentionality requires not just a behaviour pattern susceptible to an intentional description but *also* that that pattern should *not* be the result of the operation of a single mechanism. Instead, Bennett argues, we have genuine intentionality only when the intentional description reveals a unity which is not visible at the mechanistic level. Bennett thus insists on what he calls a *unity position* which posits, as a necessary condition on genuine intentional descriptions, that the description depicts as conceptually unified some set of facts which *cannot* be so unified by reference to the underlying mechanistic story. Thermostats fail the test, as a single mechanistic explanation can replace the ones citing desires to achieve certain temperatures etc. By contrast, Bennett expects that there will be no single mechanism which mediates all the behaviours which we might describe in a higher animal using a generalization such as "it is doing x because it thinks x will bring food." Why? Because the range of behaviours which might fall under this rubric includes different kinds of bodily motion and responses to different inputs. Hence "We are soberly entitled to suppose that no one mechanism explains all this behaviour" (Bennett, 1991, p.180).

Bennett thus *insists* that the true believers innards be fragmentary relative to the folk description. What Davies saw as downright inimical to the proper use of the folk talk, Bennett seems to depict as essential! My own view is that the whole approach of *counting mechanisms* is importantly misguided -not least because it depends on some very slippery notion of how to individuate mechanisms (when is it right to speak of one mechanistic route mediating an input/output pattern, and when of

two etc. I doubt if there are principled answers to such questions). The view developed in Section 1 is rather that the inner story may be fragmentary OR non-fragmentary (relative to the folk ontology) without thereby compromising the integrity of the folk talk. In short, I count the Davies/Bennett lines as orthogonal to the question of the commitments of the folk discourse.

Dennett himself does not propose any form of mechanism counting as a response to the worries we raised. But he does feel driven to concede that (Dennett, 1988, pp.542-543):

"If one gets confirmation of a much too simple mechanical explanation-this really does disconfirm the fancy intentional level account."

This is clearly the kind of intuition which Bennett sought to make precise by insisting on a multiplicity of mechanisms underlying each genuine intentional generalization. But it is hard to justify. Why should the possession of relatively simple innards unfit a being for the ranks of the True Believers? Simplicity *per se* is not a crime. And suppose it did turn out that a simple inner mechanism was mediating all of a certain sub-set of my behaviours. Why should that, in and of itself, work against an intentional/folk psychological description of those behaviours? The various moves in the debate, it seems to me, are curiously unmotivated: More an ad hoc attempt to regiment intuitions than to explain or justify them. Pure ascriptivism, just about everyone (including Dennett) agrees, won't quite do. But the shape of an alternative remains elusive.

4. *The intelligibility constraint*²

To get a better (workable) grip on the conceptual bonds linking the folk discourse and scientific studies of cognition, we need to give up a certain obsession. What has to go, I believe, is the obsession with reductive relations between types of description of complex systems. Both Davies' and Bennett's attempts to pin down the nature of the folk's inner commitments revolve around attempts to specify the necessary shape of reductive relations between individual items in the folk ontology and inner mechanisms. But a better tack, I suggest, may be to concentrate rather on the broad *properties* which the folk ascriptions assume and then to ask (of

any given account of underlying mechanisms), whether the scientific story renders intelligible the possession of such properties. If it renders the possession of a property unintelligible, and if the property seems to be close to the heart of the folk framework, we may conclude that the truth of that scientific description is indeed incompatible with the acceptability of the folk framework. Otherwise (and regardless of the neat fit, or of a lack of it, between the folk and scientific ontologies) all is well.

To get a sense of the kind of approach I have in mind, consider the case in which someone claims to have made a mistake in his deployment of a concept. Thus imagine that one day I say to you "I must admit I made a mistake when I said that all fish-meat was soft: I now recall eating some shark-meat which was quite tough." On hearing this, you might characterise me, in folk-psychological mode, as believing that I made such and such a mistake and as being prompted to recognize that mistake by recalling a certain experience. But suppose you then discover that my inner cognitive economy takes the form of a big look-up table. One implication of this is that each of my various judgements and behaviours is issued in full informational isolation from the rest. That is to say, my response today to a question such as "Are you sure that what you said about all fish-meat being soft is correct?" is in fact issued without recourse to any causal process in which a memory of my previous judgement is retrieved and a comparison of that judgement against a body of stored information is made. To learn this, I assert, is to learn that the folk-description of my behaviour (as being prompted, by recalling a previous experience, to judge that I previously judged incorrectly) is untenable. Note that it is not rendered untenable because of a failure to isolate neat internal analogues to the folk-items. Rather, it is untenable because once we know how the inner processing system operates, we can see that some very general assumptions we had been making (viz. that the 'agent' could recall previous judgements and compare such judgements against a body of stored knowledge) cannot be sustained.

Whether such a discovery would undermine the whole folk psychological enterprise would then depend on how essential these properties are to the picture of intentional agency. In the case imagined, they strike me as sufficiently central to warrant the elimination of the folk framework. Fortunately, nothing in the real scientific stories mentioned in Section 2 threatens to have any such effect. All those stories depict us as complex information processing devices capable of recall, generalization and

inference. Their only 'crime' was to fail to recapitulate the detailed taxonomy of the folk discourse. But this, I contend, is no crime at all. As long as the global properties assumed by the folk discourse can intelligibly be supposed to be possessed by beings with a given inner constitution, folk psychology is innocent.

I will give one more example of the kind of global property I have in mind, though this one may prove contentious. I am willing to assert that the folk practice of ascribing grasp of a concept assumes some kind of potential consciousness of content on the part of the beings to whom it is correctly applied. For example, it strikes me as conceptually incoherent to depict a being as, on the one hand, knowing the meaning of the concept 'dog,' and on the other hand being incapable of enjoying any conscious episode in which the idea of a dog figures. Even if I cannot persuade you of that, you might at least agree that very often, folk psychology depicts agents as consciously entertaining specific thoughts. Suppose, then, that we one day achieve a satisfying scientific theory which depicts a certain (no doubt very high-level and abstract) computational organization as a necessary condition of conscious experience.³ And suppose, in addition, that a certain recently encountered alien life-form, which we had hitherto regarded as grasping certain concepts and having a variety of conscious experiences, failed to possess the requisite inner organization. Then (unless we take this to cast doubt on the scientific theory—a live option) we should, I claim, at least conclude that parts of our folk description were wrong. And we might (depending on the extent to which we are willing to tie concept grasp to consciousness) even feel that most of our folksy descriptions of the being had been undermined. Once again, the moral is that the folk framework trades on large scale properties (consciousness, recall, and others I have not discussed such as generalization and inference) whose presence gross behaviour is simply insufficient to guarantee. Only innards whose scientific description leaves intelligible the possession of such properties are to be countenanced as fit for True Believers.

5. The bearable vagueness of believing

In attempting to dodge between the counterbalanced infelicities of the Super-Fodorian and the Pure Ascriptivist, I am exposed to a variety of

worries and criticisms. I therefore end with a brief defensive tour.

5.1 First Worry: Mental Causation

In recent years, the issue of mental causation has come to dominate the discussion of the putative virtues and vices of the folk framework. Thus Fodor (1987) clearly sees it as a major virtue of any reductive/syntactic inner story that it allows us to make easy mechanistic sense of the idea of a specific belief (or whatever) being a cause. If we are willing to give up on the hope of such straightforward reduction without thereby giving up on the folk framework, what are we to say about mental causation?

One swift, clean move is to give up on the image of beliefs as mechanistic causes and instead to focus on the (purely) explanatory virtues of the folk framework. This is the kind of move that Dennett makes in speaking of the way folk content ascriptions allow us to predict who will or will not appreciate a certain joke and so on (Dennett, 1987). The folk talk here tells us what bodies of information the subject is familiar with. This is useful information regardless of whether neat reductive analogues to specific belief contents are to be found. Something which falls far short of a detailed description of the specific inner events which enter into the push and shove of creation can nonetheless tell us a lot about the likely patterns of behaviour of other agents. We may thus give up on the individual, folk-described, beliefs as discrete causes and yet still value and exploit the intentional descriptions of agents.

Alternatively (less neatly, I concede) we may question the assimilation of causation to simple mechanistic episodes of push and shove. I remain tempted by (though I shall not attempt to defend) the idea that our understanding of causation is parasitic on our understanding of explanation and that all good explanation is, in at least some indirect sense, causal explanation. But whichever: either broaden the notion of causal to encompass global, emergent phenomena as causes, or insist that the folk-talk is explanatory, though not causal explanatory. The point is, a discourse can be powerful and valuable regardless of whether its favoured entities have neat reductive analogues which participate in the push and shove of low-level creation. (Compare: "The car crashed because of its poor cornering." Cornering is a global property not reducible to any single mechanistic fact. Instead, a car corners well or badly due to the combined influence of several internal and external factors. Yet poor

cornering, mores the pity, really can cause crashes.)

Mental causation, I conclude, is more a red herring than a fulcrum of debate, I hereby bracket it, and move on to:

5.2 Second Worry: Semi-Believers, and Demi-Semi-Believers

The characteristic items of the folk mentalistic ontology are, I suggested, in all likelihood names for highly fragmented bags of subpersonal cognitive competencies. Thus to say of someone that he grasps a given concept, or that he believes a certain proposition, is to comment on an overall (gross behavioural) competence whose internal roots may be almost (but not quite) arbitrarily fragmented. One upshot of this, which some people find uncomfortable, is that the notion of sameness of belief becomes rather fluid. Two agents (or one agent at two times) may share enough sub-competencies to count, for some purposes, as sharing the belief; yet differ with respect to enough subcompetencies to count, for other purposes, as not sharing the same belief. Likewise, if grasp of a concept is subserved by a panoply of disparate sub-personal abilities, a being may count as more or less grasping a concept according to how many such abilities it possesses. Nor need there be any neat answer to the question: if we subtract these sub-abilities, will the agent still count as grasping the concept/having the belief or whatever.

I confess to being secretly pleased with this turn of event. The macro-level folk constructs will apply to a greater or lesser extent, and in ways largely determined by the contingencies of a specific deployment of the folk discourse. In problem cases (infants, animals, brain-damaged patients -see Stich, 1983) there will indeed be no answer to the question "does the organism fall under this folk-description or not?" In respect of some sub-abilities, yes. In respect of others, no. For this purpose, yes. For that purpose, no. This strikes me as entirely intuitive. There really is no God-given answer in such cases. And (as a rule of thumb) where God fails, Philosophy and Cognitive Science had better not succeed!

5.3 Third Worry: The Fragmentation of the Global Properties.

The scientific commitments of the folk image of mind are exhausted, I claim, by some set of rather global informationprocessing properties e.g. having innards which support genuine recall and comparison, which allow

qualitative experience, etc. But what if these global properties should themselves fragment: what if, for example, we found a being who met the scientific criteria for conscious experience yet who, surprisingly, failed to meet the criteria for genuine recall -a Giant Conscious Look-up Table if you will? In such cases, I am again happy to concede that there is no good answer to the question "Is that a True Believer?" In cases where a host of features co-occurrent in, and conceptually central to, our original exemplars of a certain type come apart, we may rightly say that nothing in our previous usage determines a hard and fast answer to the question. If we then proceed to refine and alter our original conception so as to marginalize some once-central properties and to centralize others, we are engaging in a useful process of stipulative conceptual development. I am quite certain that the concept of a True Believer will undergo such change. But when it does so we should not be misled into imagining that we are literally discovering the proper extension of the original concept.

The potential fragmentation of the kind of global properties I have highlighted is thus no cause for concern. In such cases some of our folk-discourse allows us to pose questions which simply have no answers. If, for some new purpose, an answer is positively demanded, that may be a catalyst for conceptual change.

5.4 Final Worry: What If we Fail the Tests?

If the folk discourse does indeed make some assumptions about inner stories, it must be logically possible that we ourselves, seen in the naked light of scientific advance, turn out to fail all the tests. For some theorists, the mere logical possibility that *we* could turn out not to be True Believers is a compelling reduction of the attempt to allow the folk discourse to make any contact with scientific stories. For some reason, I cannot get excited about this. I cannot seem to worry about failing a test which I am completely certain I will pass! I agree that were science to one day tell me, for example, that my present outputs in fact never draw on stored knowledge about my past behaviour and experiences, that would be full-scale disaster. But in such a case, turning out not to be a True Believer would, I suggest, be about the least of my worries!

6. Conclusion: a very robust parrot

What is the relation of mind (the folk mentalistic arena of concepts, beliefs, propositions and all the rest) and cognition (the scientific arena of vectors, state spaces, languages of thought or whatever)? We sketched and rejected two extreme views. The first (SuperFodorian Realism) depicted the folk discourse as embodying a commitment to a specific, language-like inner story. The second (Pure Ascriptivism) lifted every vestige of scientific commitment from the overburdened shoulders of the folk. As an antidote to both extremes, I have tried to depict folk psychology as assuming some global, quite high-level properties which certain kinds of inner organization may demonstrably fail to underwrite. But these commitments fall far short of demanding that the specific explanatory posits of the folk discourse (concepts, propositions, attitudes) be vindicated by the discovery of corresponding inner scientific kinds. Indeed, the inner story may, consistent with underwriting the global properties, be quite fragmentary and fluid.

Paul Churchland once⁴ complained of Dennett's ascriptivism that it was akin to the insistence of Monty Python's famous pet shop owner that a certain terminally inert parrot was not dead but 'just resting'! Pure ascriptivism, Churchland argued, was just nailing the parrot to its perch. I have tried for an equally efficient but less drastic solution. The folk psychological parrot, I claim, is merely a very robust kind of animal needing only the minimum of cooperation from scientific advance to live and thrive.

George Washington University

NOTES

1. We gratefully acknowledge the permission of the Academia Sinica to reprint this paper. It appeared before in 'Mind and Cognition': collected papers from 1993 symposium on mind and cognition.
2. The use of the term 'intelligibility here may recall Cussins use of a (semi-technical) notion of 'Intelligible Connection in the development of his Construction Constraint (see Cussins (1990), pp. 374-378). But in fact the two notions are very different. Cussins notion of an 'Intelligible Connection

between levels applies just so long as the 'marching in step of the two kinds of description does not appear as a 'miraculous coincidence. But from that perspective, there exists a perfectly 'Intelligible Connection (the capital indicates Cussins semi-technical term) between e.g. the innards of the Look-up Table being and the folk level of descriptions of its behaviour. It is no coincidence -but it is not an acceptable implementation because it is UNINTELLIGIBLE (my, non-technical sense) how the folk commitments to normativity (see text) can be met by such a being.

3. For an argument against such an outcome, see Patricia Churchland (1986). For an argument in favour, see Lahev (1993).
4. In a reply to Dennett, reprinted in Churchland (1989), p. 127.

REFERENCES

- Bennett J. (1991), *Folk Psychological Explanations. The Future of Folk Psychology*, ed. J. Greenwood: Cambridge University Press.
- Churchland P. M. (1989), *The Neurocomputational Perspective*. Cambridge, MA: MIT/Bradford Books.
- Churchland P.S. (1986), *Neurophilosophy: Toward a Unified Science of the Mind-Brain*. Cambridge, MA: MIT Press.
- Clark A. (1993), *Associative Engines: Connectionism, Concepts and Representational Change*. Cambridge, MA: The MIT Press.
- Cussins A. (1990), The Connectionist Construction of Concepts, in *The Philosophy of Artificial Intelligence*, pp. 368-440, ed. M. Boden. New York: Oxford University Press.
- Davies M. (1991), "Concepts, Connectionism and the Language of Thought." *Philosophy and Connectionist Theory*, pp. 229-258, eds. W. Ramsey, S. Stich and D. Rumelhart. Hillsdale, NJ: Erlbaum.
- Dennett D. (1987), *The Intentional Stance*. Cambridge, MA: MIT Press.
- Dennett D. (1988), *Precis of the Intentional Stance (and Authors Response)*, *Behavioral and Brain Sciences* 2, pp. 495-546.
- Ellis A. and Young, A. (1988), *Human Cognitive Neuropsychology*. London: Erlbaum.
- Elman J. (1991), *Incremental Learning or the Importance of Starting Small, Technical Report 9101*, Center for Research in Language, University of California: San Diego.
- Evans G. (1982), *The Varieties of Reference*. Oxford: Oxford University Press.
- Fodor J. (1987), *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: The MIT Press.

- Humphreys G. and Riddoch, M. (1987). *To See But not to See: e Case Study of Visual Agnosia*. London: Lawrence Erlbaum.
- Lahav R. (1993), "What Neuropsychology Tells Us About Consciousness." *Philosophy of Science* **60**, pp. 67-85.
- Lycan W. (1988), Commentary on D. Dennett, the Intentional Stance, *Behavioral and Brain Sciences* **2**, pp. 518-519.
- Peacocke C. (1983), *Sense and Content: Experience, Thought and Their Relations*. Oxford: Clarendon Press.
- Ramsey W., Stich S. and Garon J. (1991), "Connectionism, Eliminativism and the Future of Folk Psychology." *Philosophy and Connectionist Theory*, pp. 199-228, eds. W. Ramsey, S. Stich and D. Rumelhart: Erlbaum.
- Ramsey W. and Stich S. (1991), "Connectionism and Three Levels of Nativism." *Philosophy and Connectionist Theory*, pp. 287-310, eds. W. Ramsey, S. Stich and D. Rumelhart. Hillsdale, NJ: Erlbaum.
- Ryle G. (1949), *The Concept of Mind*. London: Hutchinson.
- Shallice T..(1988), *From Neuropsychology to Mental Structure*. Cambridge: Cambridge University Press.
- Stich S. (1983), *From Folk Psychology to Cognitive Science: The Case Against Belief*. Cambridge, MA: MIT/Bradford Books.
- Warrington C. and McCarthy R. (1987), "Categories of Knowledge: Further Fractionations and an Attempted Integration." *Brain* **110**, pp. 1273-1296.