

Language and thought

Interdisciplinary themes

edited by

Peter Carruthers

*Professor of Philosophy and Director, Hang Seng Centre for
Cognitive Studies, University of Sheffield*

and

Jill Boucher

*Associate Professor of Human Communication Sciences,
University of Sheffield*

*Published in association with the Hang Seng Centre for Cognitive
Studies, University of Sheffield*



CAMBRIDGE
UNIVERSITY PRESS

14 Reflections on language and mind

Daniel C. Dennett

1 A seductive bad idea: central processing

A theme that emerged at the Sheffield Conference with particular force, to my way of thinking, was a new way of recognising, and then avoiding, a seductive bad idea. One of its many guises is what I have called the Cartesian Theatre, but it also appears in the roles of Central Processing, or Central Executive, or Norman and Shallice's SAS, or Fodor's non-modular central arena of belief fixation. What is wrong with this idea is not (just) that it (apparently) postulates an *anatomically* discernible central region of the brain – maximally non-peripheral, one might say – but that it supposes that there is a functionally identifiable *subsystem* (however located or even distributed in the brain) that has some all too remarkable competencies achieved by some all too remarkable means. There are many routes to it. Here is one that starts off in an excellent direction but then veers off. The mistaken fork is not explicitly endorsed by anybody that I can think of, but I daresay it has covertly influenced a lot of thinking on the topic.

- (a) One of the things we human beings do is talk to others.
- (b) Another is that we talk to ourselves – out loud.
- (c) A refinement of (b) is to talk silently to oneself, but still in the words of a natural language, often with tone of voice and timing still intact. (One can often answer such questions as this: 'Are you thinking in English or French as you work on this problem?')

So far, so good, but watch out for the next step:

- (d) A further refinement of (c) is to drop the auditory/phonemic features (indeed all the features that would tie one's act to a specific natural language) and just think to oneself in *bare propositions*.

Introspection declares that something like this does indeed occur, but we must be circumspect in how we describe this phenomenon, since the temptation is to use it to anchor a remarkable set of dubious implications, to wit:

- (e) Since propositions, like numbers, are abstract objects, they would need some vehicles of embodiment in the brain. Moreover, when one thinks in bare propositions, one's thoughts still have one feature of sentences:

logical form. (They must have logical form if we are going to explain the phenomenon of reliable deductive inference as the manipulation of these items.) So there must be a medium of representation distinct from any natural language (call it the Language of Thought, or LOT) that has this feature.

And finally:

- (f) This activity of manipulating formulae of LOT is the fundamental variety of thinking – ‘thinking proper’ or ‘real thinking’ or even (as somebody said in discussion at Sheffield) ‘where the understanding happens.’

When someone, for instance Peter Carruthers, speaks of ‘the use of peripheral modules in central processing,’ this *suggests* (without quite assuming or implying) that there is a LOT-wielding central processing system surrounding itself with a set of tools that can be put to use, on occasion, to augment the fundamental thinking powers of the LOT system. The central processing unit (CPU) of a von Neumann machine, with its requirement that all instructions be in its proprietary machine language – the only language it ‘understands’ – is perhaps the chief inspiration for this image.

But there are other ways of imagining the phenomenon of wordless thought. It can be seen as a rather exotic specimen instead of the foundation of all cognition, a type of mental event that rather rarely occurs, and depends, when it does, on the preparation of the thinker’s mind for this sort of feat by years of practice with more explicit, word-clothed or image-supported varieties of self-stimulation. Contrast the all too inviting image of a systematic central clearing-house conducting all its affairs in its inborn *lingua franca* with the image of an anarchic, *competitive* arena in which many different sorts of things happen – Grand Central Station, in which groups of visitors speaking many tongues try to find like-minded cohorts by calling out to each other, sweeping across the floor in growing crowds, waving their hands, pushing and shoving and gesturing. In this vision, most successful activities depend on enlisting large multi-modal coalitions, involving the excitation of several largish areas simultaneously, but occasionally swifter, more efficient contacts co-ordinate activity with hardly any commotion at all. These rare, hyper-sophisticated transitions occur when very sketchy images of linguistic representations serve as mnemonic triggers for ‘inferential’ processes that generate further sketchy images of linguistic representations and so forth.

It may not be words that we learn how to leave out when we engage in such wordless thought. This sort of short-cut transition can take place in any modality, any system of associations. For a trained musician, the circle of fifths imposes itself automatically and involuntarily on a fragment of heard or imagined music, just the way *modus ponens* intrudes on the per-

ceptid f a trained logician, or the offside line does on a trained soccer player. This way of trying to imagine word-free 'logical inference' makes it look rather like barefoot water-skiing – a stunt that professionals can make look easy, hardly the basic building block of successful transportation or cognition in the everyday concrete world. It is just an impressionistic image, so far, but without such images-in-advance-of-models, we tend to get sucked back into the image of the LOT-wielding CPU, as if it were the only possible way our minds could work.

We do introspect ourselves thinking, and sometimes it does seem that our thinking is wordless but 'propositional.' This is an indisputable (hetero) phenomenological fact.¹ It is bolstered by such widely acknowledged experiences as the tip-of-the-tongue phenomenon, in which we surely do have a particular content 'in mind' and are frustrated in our attempts to find the word that normally clothes it or accompanies it. Publishers have brought out humorous dictionaries of neologisms to fill in the gaps in English: 'sniglets' is proposed, I seem to recall, as a word for those handy little metal sleeves on the ends of shoelaces, and 'yinks' would name the pathetic strands of hair that some men drape carefully but ineffectively over their bald spots. We have all thought of (noticed, reflected upon) these items before – that's why these books are able to strike a humorous chord in us – so obviously we can have bare or wordless concepts in our consciousness. And so, it seems, we can think bare propositions composed of them.

This kind of thinking is a personal level activity, an intentional activity, something we *do*. (Think of the sign that admonishes 'THINK!') It is not just something that happens in our bodies. When we think thoughts of this sort, we do, it seems, *manipulate* our thoughts, and it can be difficult or easy work. The undeniable existence of this important variety of phenomena obliged Gilbert Ryle to include in *The Concept of Mind* (1949) his notoriously unsuccessful late chapter on 'The Intellect,' but also impelled him, for the rest of his career (see Ryle, 1979), to grapple with a tantalising question about *Le Penseur*, Rodin's famous chin-in-fist concentrator: What is he *doing*?

These phenomena are part of the hetero-phenomenology of consciousness, and hence form part of the explicandum of any theory. But it is *not* given to introspection, nor does it follow from any well-considered theoretical considerations, that the cognitive transitions in us that are not personal actions but do just happen in our bodies, occur in a 'propositional' medium involving the transformation or manipulation of formulae of LOT.

¹ Hetero-phenomenology is phenomenology from the third-person point of view, or in other words, the empirical, scientific study of the way it seems to individual subjects of experience. The methods and assumptions of hetero-phenomenology are explained and defended in Dennett, 1982 and 1991.

Even if the (d) phenomenon occurs quite frequently *in people like us* professional thinkers, (e) and (f) do not follow, and there are good reasons to resist them. Here is one of the most important. If we view LOT or Mentalese as the *lingua franca* of all cognition, and view it as 'automatically understood', this apparently secures a 'solution' to the problem of understanding. But in spite of declarations by Fodor and others that LOT does not itself require interpretation by the brain that speaks it, this solution-by-fiat is both unsupported and costly: it creates an artifactual problem about the 'access' of consciousness.

How? If Mentalese is the *lingua franca* of all cognition, it must be the *lingua franca* of all the *unconscious* cognition in addition to the conscious thinking. Unconscious cognitive processes are granted on all sides, and if it is conducted in Mentalese (as is commonly asserted or assumed by theorists of the LOT persuasion), getting some content translated into Mentalese cannot be sufficient for getting it into consciousness, even if it is sufficient for getting it understood. There must then be some further translation or transduction, into an even more central arena than Central Processing, into some extra system – for instance, Ned Block's (1991) postulated consciousness module. Beyond understanding lies conscious appreciation, according to this image, and it needs a place to happen in.² This is what I have called the *Myth of Double Transduction*, and I have criticised it elsewhere (Dennett, 1996a), and will not pursue it further here.

2 A better idea: contention scheduling all the way up

The person – the star of the personal level – who does the thinking is not and cannot be a central subsystem. The person must be the supersystem composed of the subsystems, but it is hard to understand how this could be so.

If 'Central Processing' were understood to be just the name of a central-ish arena in which competitions are held, with no knowing supervisor, it would not be an objectionable term. If there is no central agent doing things – and hence no problem of *this agent's* 'access' to anything, then the only access that needs accounting for is the access that one subsystem or another has to the fruits of another subsystem. (We could acquiesce in the jargon and call these modules or quasi-modules, but unlike Fodorian modules, they do not pour their output into Central Processing, there to be 'accessed' by . . . *Moi*.) We have to get rid of the Central Understander, the Central

² If we follow Jackendoff's idea (1987), this would not be the *central* summit but a sort of ring surrounding that summit – a tempting idea, but not, I think, one to run with. (I'd rather get the benefits of Jackendoff's vision by other routes, other images, but that's a topic for another occasion.)

Meaning the Central Thinker. And to accomplish this, the thoughts are going to have to think themselves somehow.

How? By 'contention scheduling' all the way up. (I am particularly grateful to Josef Perner for re-introducing Norman and Shallice's model into the discussion at Sheffield.) Norman and Shallice (1980) contrasted 'automatic' contention scheduling – an unsupervised competition between independent modules that handles the routine bulk of unconscious conflict resolutions – with the labours of the Supervisory Attentional System (SAS), an ominously wise overseer-homunculus who handles the hard cases in the workshop of consciousness.

Shallice's 1988 book has some valuable insights on how to soften this contrast, but I propose to reject it altogether. How could the Supervisory Attentional System do its work? To the Norman and Shallice model of low-level contention scheduling we have to add layers (layers upon layers) of *further* contention scheduling, with suppression, coalition, subversion, and so forth of higher-level contestants.

How could the SAS suppress the habits determined by lower-level contention scheduling? A homunculus with big strong arms? No. These habits have to be resisted by counter-habits, things of basically the same sort: further contestants, in other words. What are these contestants? We might as well call them 'thoughts' made out of 'concepts,' but we mustn't understand these in the traditional way. You don't move your thoughts around; your thoughts move you around. You don't make your thoughts; your thoughts make you.

Several participants at Sheffield expressed the basic idea: 'cognition is deeply involved in dialogical activity', 'cognitive and communicative factors merge'. How could thinking be anything other than communicating? The bad idea supposes that propositions get 'grasped' by the mind (as Frege said) when they get expressed in Mentalese. But how could a brain's central system merely writing Mentalese in itself count as thinking? How could that do any work, how could that guarantee understanding? Only by enabling something. What? Enabling the multitudinous items of information that are variously distributed and embedded around in the brain to influence each other, and ongoing processes, so that new, better informational structures get *built* (usually temporarily) and then further manipulated. But the manipulanda have to manipulate themselves.

For many years, Douglas Hofstadter has spoken of 'active symbols', an epithet which surely points in the right direction, since it moves us away from the idea of inert, passive symbols that just lie there until manipulated by some knowing symbol-user. The germ of truth in the idea of using the von Neumann machine CPU as our model of Central Processing is the fact

that each member of the CPU's instruction set is a sort – a rigid sort – of active symbol. Whenever it appears in the instruction register, its mere tokening there makes something happen; it triggers a specific activity – adding or multiplying or checking the sign of the number in the accumulator, or whatever – because its defining shape *turns on* a particular bit of special-purpose circuitry. These symbols are all imperatives, the 'understanding' of which consists in automatic obedience. There may be hundreds of specialised circuits or just a few (as in the recent return to Reduced Instruction Set Computers); in either case the vocabulary is in tight correspondence to the hardware (or micro-code). Unless Mentalese is considered to be restricted to such imperatives of internal operation (and I have never encountered so much as a hint of this in philosophers' discussions of it), there is simply no analogy at all between Mentalese and machine language. So the promise of 'automatic understanding' of Mentalese or LOT is an empty one.

3 Making tools to think with

Now we need to work out how active symbols might come into existence. Andy Clark and Annette Karmiloff-Smith (1994) contrast embedded concepts with those that are disembedded via re-representation, and Josef Perner (this volume, chapter 13) speaks of explicit concepts. How does this re-representation or explicitation multiply the powers of a concept? Part of the answer may come from reconsidering Köhler's classic set of experiments with problem-solving apes (1925).

Contrary to popular misunderstanding, Köhler's apes did not just sit and think up the solutions. They had to have many hours of exposure to the relevant props – the boxes and sticks, for instance – and they engaged in much manipulation of these items. Those apes that discovered the solutions – some never did – accomplished it with the aid of many hours of trial and error manipulating. Now were they thinking when they were fussing about in their cages? What were they manipulating? Boxes and sticks. It is all too tempting to suppose that their external, visible manipulations were accompanied by, and driven by, internal, covert manipulations – but succumbing to this temptation is losing the main chance. What they were attending to, manipulating and turning over and rearranging were boxes and sticks, not thoughts.

We *Homo sapiens* engage in much similar behaviour, moving things around in the world. For instance, most Scrabble players would be seriously handicapped if they were prevented from sliding the little tiles around on their little shelf. We write words on index cards and slips of paper, doodle and diagram and sketch. These concrete activities are crutches for thinking,

and as we get practised at these activities, we can internalise these manipulations.

Can you alphabetise the words in this sentence in your head?

Yes, probably, but what you do is not easy. It is not easy because, as Keith Frankish stressed in Sheffield, you have to *keep track of* things. To perform this stunt, you need to use visual imagery and auditory rehearsal. You work to make each component as vivid as possible, to help you do the required tracking. If we were to expand the task just a little bit by making the words hard to distinguish visually, or by lengthening the sentence, you would find it impossible to alphabetise in your head – and this very sentence is such an instance. Try it.

As Barry Smith noted at Sheffield, the metaphysical claim that thought depends on language can be traded in for something much more interesting: the scientific hypothesis that there is a deep natural necessity for thought – our kind of thought – to involve language. As he pointed out, ‘Thoughts are not the sort of thing you can grasp like an apple (*pace* Frege)’ – but words and sentences *are* things you can grasp like an apple. They are *manipulanda*, like Köhler’s boxes and Scrabble tiles.

For the moment, set grammar (and logical form) aside and think of words just as isolated objects – images or labels, not parts of sentences obeying the relevant rules. This is not a natural act for a philosopher of language, so impressed have we all been with the combinatorial and analytic power of syntax. But you don’t have to tie the good idea of inference to logical form. Some of us do on occasion think whole arguments in our heads – because we’ve taken and passed logic class. But it is obviously true that most people never engage in *explicit* non-enthymematic formal reasoning; whether people covertly or unconsciously do their thinking by symbol manipulation ought to be a matter of controversy at best. Johnson-Laird and others have urged us not to take this as our model for rational, useful thought. What other models do we have? Johnson-Laird’s account of mental models (1983) is one, but I’ll mention another: Hofstadter’s model of analogical thinking (Mitchell, 1993; Hofstadter, 1995; French, 1995).

In Hofstadter and Mitchell’s Copycat programme, a simple analogy-finding game is played, using the alphabet as the toy world in which all perception and action happens. Suppose it is my turn: I make a move by altering some alphabetic string. You must respond by altering your string (which is given in advance) ‘in the same way’. Thus if my move is to turn ‘abc’ into ‘abd’ it is pretty obvious that your response should be to turn ‘pqr’ into ‘pqs’ (advancing the last or rightmost symbol in an alphabetic sequence by one). But what if your sequence was not ‘pqr’ but ‘bbccdd’? Change it to ‘bbccce’ probably, treating the repetition of the letters as cre-

ating twin-elements to be treated as one. The ‘probably’ in the previous sentence is key. There are no correct or incorrect answers in Copycat, only better or worse, more or less satisfying or elegant or deep responses.

Consider if my move is to change ‘fgh’ to ‘ffggghhhh’. What is your best move if your initial string is ‘cba’? It might be ‘ccbbbaaaa’ or you might think that alphabetical order is more important than left-right order, and go instead with ‘cccbbbaa’. The Copycat programme plays this game. You give it your move (and its initial string) and it figures out a response. In the course of its deliberations, different ‘concepts’ are *built*, out of the existing resources. They are built on the fly, and allowed to evaporate when no longer needed. (A point also made by Sperber and Wilson, this volume; chapter 9.)

The Copycat programme does well on many Copycat problems, but there are classes of Copycat problems we human beings find relatively easy that are beyond it. Consider, for instance, your best response in the following case. I change ‘abcdjf’ to ‘abcdef’, and your string is ‘pwppp’. Isn’t it obvious that to do the same thing to your string you should turn it into ‘ppppp’? We English speakers have a word (and concept) for this move: ‘repair’. I fixed my defective alphabetic string, so you must fix your defective five-of-a-kind string. The concept of *repair* is a tool in your kit, but it simply cannot be built from the existing resources (to date) in Copycat. Similarly, as Andy Clark (this volume, chapter 8) demonstrates in his discussion of Thompson *et al.* (forthcoming), chimps can have their repertoire of *usable* concepts enlarged by adding a symbol for ‘sameness’. This discrimination becomes available or usable by the chimps only when they are given a crutch: a token or symbol that they can lean on to help them track the pattern. The thought processes exhibited here (by Copycat, by the chimps) are familiar human thought processes, and they are not logical arguments; they are (roughly) processes of competitive concept-building. In order to engage in these processes, however, one must be able to keep track of the building blocks, and tracking and recognition are not for free. Our concepts are clothed in re-identifiable words for the same reason the players on a sports team are clothed in uniforms of the same *familiar* colour: so that they can keep track of each other better (so that they can find each other readily in the Grand Central Station of the brain).

4 What we can do with these tools that no other animals can do

If we are enabled to know things other animals don’t know about their own minds, this has to have some payoff. There must be things we can do that they can’t do. There are: lots. As Carruthers (this volume, chapter 5) says, the point of human consciousness is to make various mental contents

repeatedly available for further processing. Consciousness enables us to *say* (to others) what we're doing. Less trivially, it enables us to *say to ourselves* what we're doing. And when we do this, we find we can (often) understand what we're saying! We can then use this little bit of extra leverage, leverage provided by our new recursive tool, to learn how to do things better. As Barry Smith noted at Sheffield, knowing you have a belief gains you leverage you don't have by just having that belief.

Consider a familiar human activity that we rely on in many problem-solving circumstances: we 'query the belief box', as somebody put it in Sheffield. We ask ourselves explicit questions. This practice has no readily imaginable counterpart in non-linguistic animals, but what does it gain us, if anything? I think the answer can be seen in Plato's analogy, in the *Theaetetus*, between human memory and an aviary:

SOCRATES Now consider whether knowledge is a thing you can possess in that way without having it about you, like a man who has caught some wild birds – pigeons or what not – and keeps them in an aviary for them at home. In a sense, of course, we might say that he 'has' them all the time inasmuch as he possesses them, mightn't we?

THEAETETUS Yes.

SOCRATES But in another sense he 'has' none of them, though he has got control of them, now that he has made them captive in an enclosure of his own; he can take and have hold of them whenever he likes by catching any bird he chooses, and let them go again; and it is open to him to do that as often as he pleases.

Possession is good, but not much use unless you have the ability to get the right bird to come when you need it. How do we do it? By means of technology. We build elaborate systems of mnemonic association – pointers, labels, chutes and ladders, hooks and chains. We refine our resources by incessant rehearsal and tinkering, turning our brains (and all the associated peripheral gear we acquire) into a huge structured network of competencies. In our own case, the principle components of this technology for brain-manipulation are words, and no evidence yet unearthed shows that any other animal is capable of doing anything like what we do with our words.

Have you ever danced with a movie star? Do you know where to buy live eels? Could you climb five flights of stairs carrying a bicycle and a cello? These are questions the answers to which were probably not already formulated and handily stored in your brain, and yet they are readily answered reliably by most people. How do we do it? By engaging in relatively effortless and automatic 'reasoning'. (See Powers, 1987 for valuable reflections on these processes.) In the first case, if no recollection of the presumably memorable event is provoked by considering the question, you conclude that the answer is No. The second question initiates a swift survey (pet stores? fancy restaurants? fish markets or live bait dealers?), and the

third provokes some mental imagery which ‘automatically’ poses the relevant further questions, which, when posed to the ‘belief box’, yield their answers. That is how we get the right birds to come – by asking ourselves questions (as Socrates noted) and discovering that we know the answers.

5 Consciousness: ‘access’ for whom?

Perner suggested at Sheffield that ‘predicative representation is necessary for consciousness’ a theme also expressed with variations by Carruthers and Smith. I suggest that such explicit predicative representation is typically sufficient for consciousness, but not necessary. What *is* necessary? Just the sort of cerebral dominance I have analogised to *fame*: Consciousness is more like fame than television (1996b). Contents ‘enter consciousness’ (a very misleading way of speaking) by being temporary winners of the competitions, persisting in the cerebral arena, and hence having more and more influence, more and more staying power (in memory – which is not a separate system or box). As Michael Holderness has aptly observed, the winners get to write history – indeed, that’s what winning is, in the brain. One very good way of achieving cerebral celebrity is to form lots of coalitions with words and other labels. All this has to *happen* in the central arena, in ‘central processing’, but not under the direction of anything like a subsystem. The person is the Virtual Governor, not a real governor; the person is the *effect* of all the processes, not their cause.

A common reaction to this suggestion about human consciousness is frank bewilderment, expressed more or less as follows: ‘Suppose all these strange competitive processes are going on in my brain, and suppose that, as you say, the conscious processes are simply those that win the competitions. How does *that* make them conscious? What happens next to them that makes it true that *I* know about them? For after all, it is *my* consciousness, as I know it from the first-person point of view, that needs explaining!’ Such questions betray a deep confusion, for they presuppose that what *you* are is something *else*, some Cartesian *res cogitans* in addition to all this brain-and-body activity. What you are, however, just *is* this organisation of all the competitive activity between a host of competencies that your body has developed. You ‘automatically’ know about these things going on in your body, because if you didn’t, it wouldn’t be your body!³

³ Barry Smith noted at Sheffield that ‘there is a way our minds are known to us that is not available to animal minds’. I agree, but I am inclined to disagree with his softening of this striking claim: ‘There is no reason to deny them an inner life.’ There is indeed a reason: They aren’t *first persons* in the way we are. They don’t have to be, so anything they have in the way of an inner life must be so dimensionally thin, so impoverished to the vanishing point as hardly to count as an inner life at all.

acts and events you can tell us about, and the reasons for them, are yours because you made them – and because they made you. What you are is that agent whose life you can tell about. You can tell us, and you can tell yourself. The process of self-description begins in earliest childhood, and includes a good deal of fantasy from the outset. (Think of Snoopy in the *Peanuts* cartoon, sitting on his doghouse and thinking ‘Here’s the World War I ace, flying into battle. . .’) It continues through life. (Think of the café waiter in Jean Paul Sartre’s discussion of ‘bad faith’ in *Being and Nothingness* (1943), who is all wrapped up in learning how to live up to his self-description as a waiter.) It is what *we* do. It is what *we* are.⁴

Several speakers in Sheffield drew attention to cognitive abilities that are particularly human, and that thus raise the question of whether they can be shared even in rudimentary form by animals without language. Perner, for instance, drew attention to Norman and Shallice’s list of the five specialities of the SAS: planning, troubleshooting, dealing with novelty, dealing with danger, and overcoming habits. There are plenty of familiar anecdotes proclaiming that birds and mammals – at least – exhibit these talents on occasion, but the very fact that these anecdotes get retold shows that they recount remarkable and impressive cases. Just how good are chimpanzees, really, at these five accomplishments, for instance? In addition to the anecdotes of glory, there is evidence, both experimental and anecdotal, of their widespread failure to rise to challenges of this sort. It is not easy to design experiments that test in both language-free and unequivocal fashion for such skills as troubleshooting or dealing with novelty, but the design project promises to repay us for our efforts in two ways. First, even when we are thwarted in our attempt to design a suitable experiment, the obstacles encountered may illuminate the role that language plays in our own case, and second, when we succeed, our experiments promise to clarify further the limits of non-linguistic thinking by other species.

⁴ Parts of the preceding paragraphs are drawn, with slight revisions, from Dennett (1996a).

References

- Acredolo, L. P. and Goodwyn, S. W. (1988). Symbolic gesturing in normal infants. *Child Development*, 59, 450–66.
- (1990). The spontaneous development of symbolic gestures. In V. Volterra and C. Erting (eds.), *From Language to Gesture in Hearing and Deaf Children*. New York: Springer-Verlag, pp. 68–78.
- Allen, C. and Bekoff, M. (1997). *Species of Mind: The Philosophy and Biology of Cognitive Ethology*. Cambridge, MA: MIT Press.
- Anderson, J. R. (1983). *The Architecture of Cognition*. Cambridge, Mass: Harvard University Press.
- Anderson, M. (1992). *Intelligence and Development: A Cognitive Theory*. Oxford: Blackwell.
- Andrews, J., Livingston, K. and Harnad, S. (submitted). Categorical perception effects induced by category learning.
- Antony, L. (1989). Anomalous monism and the problem of explanatory force. *Philosophical Review*, 98, 153–87.
- (1991). The causal relevance of the mental. *Mind and Language*, 6, 295–327.
- Armstrong, D. (1980). *The Nature of Mind and other Essays*. Ithaca: Cornell University Press.
- Atran, S. (1990). *Cognitive Foundations of Natural History*. Cambridge: Cambridge University Press.
- Atkinson, R. L., Atkinson, R. C., Smith, E., Bem, D, and Hilgard, E. (1990). *Introduction to Psychology*. San Diego: Harcourt Brace Jovanovich.
- Avramides, A. (1989). *Meaning and Mind: An Examination of a Gricean Account of Language*. Cambridge Mass.: MIT Press.
- Baddeley, A. (1986). *Working Memory*. Oxford: Oxford University Press.
- (1988). *Human Memory*. Hillsdale, NJ: Lawrence Erlbaum.
- (1996). Exploring the central executive. *The Quarterly Journal of Experimental Psychology*, 49A (1), 5–28.
- Baddeley, A. and Hitch, G. J. (1974). Working memory. In G. Bower (ed.), *The Psychology of Learning and Motivation*, vol. 8, pp. 47–90. New York: Academic Press.
- Baier, A. (1979). Mind and change of mind. *Midwest Studies in Philosophy*, 4, 157–76.
- Barkow, J., Cosmides, L. and Tooby, J. (eds.), (1992). *The Adapted Mind*. Oxford: Oxford University Press.
- Bar-On, D. (1995). ‘Meaning’ Reconstructed: Grice and the naturalizing of semantics. *Pacific Philosophical Quarterly*, 76, 83–116.