

COMMENTARY

Cognitive Maps, Time and Causality

JOHN O'KEEFE

*Department of Anatomy and Developmental Biology,
University College, London, WC1E 6BT*

FIFTEEN YEARS AGO, Nadel and I proposed a theory of navigation and spatial representation in animals which broached several topics of interest to students of the philosophy of space. We addressed many of these in the introductory chapter to that book (O'Keefe and Nadel, 1978). In the intervening years the theory has been substantially developed and extended and I have tried to relate some of these changes to the philosophical issues (O'Keefe, 1985; 1993).

John Campbell, in his paper, has discussed several of these ideas and suggested alternatives. In my response I should like to concentrate on some of his points, either to clarify my position or to show how it differs from his. In particular, I should like to address the questions of

1 the *objectivity* of the cognitive mapping system — to what extent it can be said to give a representation of the world independent of the organism;

2 the role of objects in the construction and use of the mapping system — a Kantian position would commit one to the notion that the spatial framework is ontologically prior to the concept of an object;

3 the roles of causality and time in the construction of an objective spatial world.

In order to lay the groundwork for the discussion of these points, I will sketch out the basic theory and give some flavour of one computational model which I have produced to place some of the predictions of the theory on a quantitative basis. Many of the details of the argu-

ments and the evidence on which they are based will be omitted from the present paper but can be found in the works listed in the bibliography. Although this will allow me to address some of the issues in a general way, I believe, in agreement with John Campbell, that some of the questions which arise at the philosophical level may be answerable only when detailed scientific understanding of the phenomena is available. Conversely, our scientific approach to these phenomena needs to be guided and motivated by the philosophical concerns.

1. The cognitive map theory

Nadel and I proposed that there are three basic strategies which animals (including humans) can use to find their way around the world, and that these different strategies are independent and are subserved by different parts of the brain. These strategies are called guidances, orientations and maps. Guidance strategies are based on the selection of an environmental feature and the maximization or maintenance of some aspect of this feature. For example, a landmark such as a large mountain could be chosen as a guide and behaviour selected which would increase the size or brightness of the retinal image of this environmental feature. An alternative strategy, the orientation strategy, involves a more specific association between the stimulus and response. When a particular stimulus is received, a behavioural response is emitted. For example, an animal might learn, as behaviourists such as Clark Hull suggested, to turn left at the corner or to dig at the base of the large elm tree. Strung together, guidances and routes comprise a procedure for getting from one part of an environment to another — in short, they make up the ingredients of a route.

In opposition to these two types of behaviourist or *dispositional* strategy the theory postulates a third: the construction and use of a map-like *representation* of the environment. Cognitive maps, the representations used by this system, consist of a set of places and a mechanism for linking these place representations together in terms of the distance and direction between them. These maps are located in a specific part of the brain called the *hippocampus*, a primitive form of cortex.

The original evidence for the theory came from the finding that cells in this part of the brain were active when the animal went to a particular location in an environment and not elsewhere or during

other behaviours. Subsequent work has identified cells in a neighbouring part of the brain which signal the direction in which the animal is pointing in an environment, irrespective of its location in that environment. In addition, there is evidence that information about the speed with which an animal moves during movements is also available to this region of the brain. We have, then, the three ingredients required to build a map-like representation of an environment and to use it for navigation: places, directions and distances (the integral of velocity). Further, there is now extensive evidence that damage to this part of the brain severely disrupts an animal's ability to use the mapping strategy although there is no effect on route strategies. For example, such animals cannot learn to find a platform in a swimming pool if it is hidden below the surface of the water but have no trouble when it is visible.

In a recent extension to the cognitive map theory (O'Keefe, 1991), I have suggested one way in which an environment could be represented within a network of brain cells. On this model, the shape of an environment could be captured by the computation of two (or more) parameters about the configuration of features in that environment. These parameters would uniquely identify that environment and distinguish it from others, and, if they fulfilled certain criteria such as orthogonality, could function as the bases for a spatial framework. Two such parameters are the centroid, the geometrical centre of the environment, and the slope or elongatedness of the environment (see Figure 1). If these could be computed for each environment, they could form the bases for a polar coordinate system, with the centroid acting as the origin and the slope acting as a reference direction against which angles could be measured. Within this polar coordinate system, items could be located relative to these axes independent of the location of the observer. Further, the neural network which captures this framework is probably organized so that, with learning, the groups of neural elements which represent environmental features come to support each other in such a way that the removal or obscuring of some of them has only minor effects on the overall spatial framework. We can suppose, for example, that cues which are close to each other would mutually support each other in such a way that they could substitute for each other, while those farther away would remain independent in their contribution to the representation of the environment. As we shall see, this ability to associate disparate cues may provide the basis for the concept of an object.

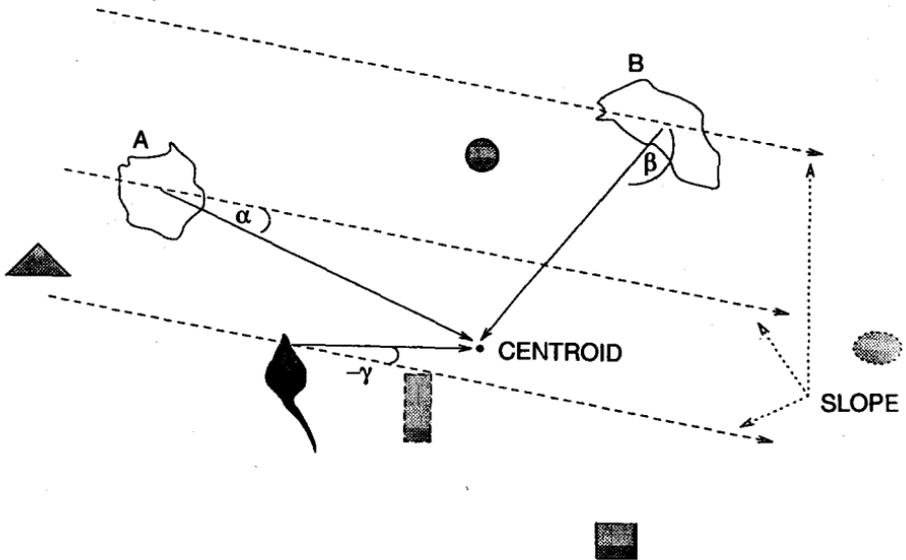


Figure 1. Representation of the centroid and slope of an environment. The environment consists of five cues which may be features. The centroid is the geometric centre of the cues while the slope is a direction which is based on the asymmetry of the cue distribution. Both are characteristic of the particular environment and can be calculated from any location in the environment. An animal such as a rat can locate itself in an environment using the centroid and the slope as the bases of a polar coordinate system. The centroid acts as the origin and the slope as the 0° direction. The rat locates itself in terms of a vector whose length is the distance to the centroid and which makes an angle (γ) with the slope direction. Other places in the environment (A & B) can be similarly located. The animal can use the mapping system to remember the locations of items such as food or water and can compute a trajectory from its current location to one of these desired goal locations by vector addition or subtraction. Detail discussion of the computational machinery involved and plausible suggestions for its neural instantiation can be found in O'Keefe (1991).

2. What philosophical conclusions can we draw from the model as exemplified in an animal such as the rat?

The most obvious conclusion is that the spatial framework is a representation created by the brain, that the existence of a corresponding framework in the physical world is an open question: all that is required is that there be a reasonable correspondence between the structure of the physical world and the elements of the spatial representation system. For example, the physical world might be something like a gigantic virtual reality machine in which stimuli were fed to the senses in an order dependent on the pattern of an organism's movements, but

the 'sense' of space was not part of the input data. In this scenario there need be no places or even spaces as such in the physical world: they are a creation of the mind. Here we appear to have neurological support for the Kantian claim that space is a part of the mental machinery for knowing the world, and may not be a property of the *Ding an sich*. (See O'Keefe, 1993 for further discussion.)

The second conclusion relates to the Kantian claim that space is ontologically prior to the concept of objects, and provides the framework within which objects can be individuated and identified. Here, my argument is as follows. There is no strong evidence that animals such as rats have the concept of an object (see below). Further, the model of the cognitive map which is being proposed is one which operates quite well on features where these are conceived as changes in the stimulus array within a single sensory modality (e.g. the items labelled A, B etc. in Figure. 1 might refer to a patch of light, a tone of a particular pitch, a pungent smell). Therefore it would appear that the concept of space can exist independently from the concept of objects. Here, John Campbell and I are in complete agreement. This is not enough to support the Kantian position that space is prior to objects. It might simply mean that they are independent, in which case we might expect to find animals which had the concept of an object but no concept of space. I shall argue below that the development of the concept of an object may be facilitated by the cognitive mapping system, but leave open the possibility that it might develop independently of notions about space or that once developed it might be elaborated in non-spatial ways.

The third conclusion I wish to draw is that the mapping system in animals such as the rat generates a representation of the environment which can be appropriately termed 'objective'. By objective, I mean a representation which is conceived as independent of the organism in terms of either its actions or its purposes. The environment is conceived as a stable entity through which the animal moves, and which exists independently of the animal's motivations, goals or desires. Here, John Campbell and I clearly disagree. He wishes to deny this premise for two reasons. First, he notes that the computational model outlined above relies on vector computations made from the cue array as perceived from the current location of the animal. Furthermore, part of the computation relies on the use of abstract representations of the animal's movements in terms of speed or distance moved. How can an

objective notion of space be derived from a system which relies on egocentric sources of information, he asks?

Campbell's second objection to calling the cognitive map of the rat an objective allocentric spatial system is his belief that the representation that it generates is ineluctably impregnated with the animal's purposes and goals. He accepts that the space may be built on the basis of features but its use in navigation is related (restricted) to achieving the animal's goals and purposes. The rat uses the map to get to places which are believed to contain desirable things such as a nest or food, and these (according to Campbell) are probably conceived of, by the animal, as objects.

I have conceded the point that, in the rat, the cognitive mapping system generates a representation which in some sense is fixed to a particular point of view, that of the animal's current physical location (O'Keefe, 1993). Changes in the focus are achieved by an internal coupling between a corollary discharge from the voluntary movement system to the mapping system which predicts the direction and distance of shift in the map which will result from the current movement. This tight linkage between actual movement and shift in the map does not appear to be an intrinsic limitation of the system, however, and I have speculated that in other, more cognitively developed, animals the point of view could be uncoupled from the location of the animal's body and imagined to move to other parts of the environment. This would be accomplished by uncoupling the movement of the focus within the map from the animal's actual movements with the development of a mechanism for the suppression of the motor signal to the musculature. These intended or anticipated movements would now afford the animal a means for representing an environment from all possible perspectives or — what is equivalent — from no particular perspective.

My response to John Campbell's second objection is that he is conflating two uses of behaviour and two types of motivation. The cognitive map theory specifies that maps are built not to satisfy biological needs which the animal currently entertains, but on the basis of curiosity, the motivation to know an environment, to know where one is. Here, nature appears to have taken an extraordinary gamble: the acquisition of information for its own sake may confer biological advantage. When an animal first encounters an unknown environment, it searches for a representation of that environment but fails to find one. This mismatch triggers exploration — a behaviour which is designed to acquire the information for building a map. There is no scientific

evidence on the exact mode of control of the motor system during exploration, but one can speculate that, *pari passu*, these movements are such as to enable the animal to acquire the maximum information with the minimum effort or in the minimum time. This might involve moving from known to unknown areas or at certain angles to unknown features. The behaviour is driven by the requirements of the spatial information-gathering system. Once in possession of a map of an environment, the animal can put it to all manner of purposes. When hungry, it can navigate from any point in a known environment to food locations; it can find its nest or avoid known danger spots. It can add additional information as it becomes available. When a specific need arises, it already has the information available to navigate to the appropriate goal despite the fact that the map was not specifically constructed with that need in mind. Aside from the advantage of using a single representation to solve many potential needs, this cognitive 'information for its own sake' strategy has the advantage that the learning can be spread over long periods of time when the animal is relatively sated, and need not be confined to particular periods of the day or to those periods when the need has already arisen. The map can be used to satisfy needs, but it is neither built on the basis of those needs nor is it permeated by them. It is entirely possible for animals to have maps of environments which contain no biologically significant objects and in which they have never experienced needs or desires.

One possibility which must be considered is that animals such as the rat do not have the concept of an object at all. John Campbell appears to believe that they do but the evidence is not compelling. It would seem entirely possible that most or all of the behaviour of these animals is mediated by unimodal features acting independently of each other and not tied to the concept of a multimodal object. The rat need not have the concept of a nest, or offspring, or even of another rat in order to discharge its biological duties effectively. Aside from the cognitive mapping system, many of its biologically motivated behaviours can be understood as a sequence of independent stimulus-response reflexes, each reflex contributing to the situation which elicits the next one. Even within the cognitive mapping system, the animal might want to go to a place but expect to find a simple stimulus rather than an object when it gets there. In the next section I shall discuss how the existence of the mapping system might provide a prerequisite for the construction of the concept of an object.

3. What is the relationship of the allocentric mapping system to the concepts of object and causality?

If it be conceded that the mapping system is phylogenetically, and perhaps ontogenetically, prior to the concept of objects, can the case be made that it is also ontologically prior, as Kant assumed? Before I discuss this, I must give my definition of the core element contained in the concept of an object, which differs from that of John Campbell. He wishes to apply the concept of object to a set of interacting dispositions and to associate with these the notion of internal causality: the idea that at least one of the causes of the present object is its past history. My notion of object is much simpler and more easily related to the spatial system. Recall that the mapping system is particularly interested in the spatial relations between features. Those features which maintain spatial coherence from exposure to exposure are incorporated into, and provide the foundations for, the map. Furthermore, the spatial separation between features is important: those which are close together or in the same location are linked in the map in such a way that they mutually support each other and can substitute for each other. I should like to suggest that this capacity to associate features and to substitute one for the other forms the basis for the primitive notion of an object. The elemental definition of an object then is a bundle of features which occupies the same location in the map and which coheres under movement of the observer or of the features themselves. In the former case the feature bundle can be incorporated as a landmark into the map, but in the latter it can not. It can, however, still be treated as an object, and may, for example, be used to predict the future configuration of the stimuli in the map. One of the essential developments here might be a mechanism for sharply delineating the boundaries of features in order to better locate them within the mapping system and to facilitate the aligning of the borders of closely overlapping features: for example, a visual configuration which closely, but not exactly, matches a tactile one might be (mis)perceived as doing so. Note that there is no need for the concept of causality here. This is just as well since I am a little worried about John Campbell's use of the concept of *internal* causality to support the concept of an object, because it seems to incorporate the notion of a time over which an object can endure and the influence of the object on itself across this time. But what is this time, does it exist in the physical world, and if so, how is it perceived or conceived by the organism? Since Campbell

subsequently employs a notion of time to elaborate the cognitive maps into narratives, he should say what the relationship between these two concepts of time is. The concept of object that I have outlined above does not rely on a concept of time, but on the simpler notion of the same configuration of features in the same place or in a predictable place.

On this view, the concept of an object is not inextricably bound to the concept of matter, but rather the latter is a postulate that is invoked *to explain* the constant conjunction of features in the same place. The light and sound occur in the same place because there is some underlying substance to which each attaches, from which each emanates. But would not an insubstantial three-dimensional glowing buzz be accorded the status of an object? Is the notion of a non-physical object necessarily a contradiction?

My concept of causality is also different from John Campbell's. I should like to restrict causality to the connectedness between environmental features which is not captured by their role in the spatial framework or in the concept of objects. Recall that features may remain spatially coherent to each other and they may be near to or far away from each other. Features that are distant from each other and yet coherent are preferentially used to construct the spatial framework while features which are close enough to overlap, whether they cohere with the rest of the spatial features or not, are treated as the basis for objects. Causation on this analysis might be the concept applied to the third type of regularity between features: one in which a spatial change in one feature (a primitive event) occurs at the same time as, or shortly after, a spatial change in another. This temporal conjunction between feature events is the basis for assigning a connection between them. It is what is left over when space, time and objects are accounted for. Notice that this notion of causality does not involve the concept of objects but it does require a concept of time, since for one event to cause another it must precede it. This concept of time, however, may be a fairly limited one, it being necessary only to assign an order to the events and to confine the time lapsed between them to a reasonably short period, of the order of 100 milliseconds or less according to the experiments of Michotte. More sophisticated notions of cause arise from attempts to rule out the possibility that the two observed events are linked, not by a direct effect of one on the other, but because they are both 'caused' by a third, hidden event. This can be ruled out

where it is possible to cause the first event directly and observe the usual effect on the second.

4. The cognitive map and linear time

John Campbell imports the ideas of objects and time into the spatial framework in his quest for an 'objective' spatial system. As we have seen, it is possible that the concept of objects can only be extracted from the feature array with the help of the spatial system. As he says, the addition of a temporal dimension to the mapping system provides a spatio-temporal framework which can be used for the generation of the concept of episodes or narratives. In our 1978 book, Nadel and I could not find any evidence that animals such as the rat had a temporal component to their spatial map. And as far as I am aware, there is still no evidence to support this idea. On the other hand, there was considerable evidence that humans *do* have a spatio-temporal system, and we suggested that this system provided a memory system for episodes and their linguistic counterpart narratives. The notion of time which is needed here, as John Campbell notes, is not the cyclical time provided by circadian rhythms, but a concept of linear time. This is a time which is unidirectional and which can be used to fix each experience uniquely.

Although the simple cyclic circadian clocks which time the day of all living organisms are clearly not sufficient to provide this time sense, there may be a relationship. It is known that periods longer than a single day can be timed using the circadian clock. For example, the four-day oestrous cycle of the female rat is dependent on the integrity of the one-day circadian rhythm. Even more interesting is the fact that the annual migration of some birds has been tied to the interaction of two circadian clocks that run at slightly different periods. This difference results in a slow predictable shift of the phase relationship between the clocks. The overall phase pattern repeats once a year and can be used to time annual events. For example, when the two clocks are in phase, the birds would know that it was spring and would prepare to migrate north to their summer breeding ground. Two clocks 180° out of phase would signal autumn and time to migrate south to the wintering areas. If one extended this system to three or four daily rhythms with slightly differing periods, it would be possible to compute daily, yearly and lifelong signals which together would provide the

signal for time-stamping each occurrence of map activation with a linear, non-repeating time.

Finally, I am in agreement with John Campbell that the existence of a spatio-temporal mapping system provides the basis for a concept of a self. On this view, the self is the representation of one's own body within such a system (see O'Keefe, 1993 for further discussion).

Bibliography

- Allison, H. E. 1983: *Kant's Transcendental Idealism*. New Haven: Yale University Press.
- Baron-Cohen, S., Leslie, A. M. and Frith, U. 1985: Does the autistic child have a 'theory of mind'? *Cognition*, 21: 37-46.
- Bennett, J. 1966: *Kant's Analytic*. Cambridge: Cambridge University Press.
- Bisiach, E., Berti, A. and Vallar, G. 1985: Analogical and logical disorders underlying unilateral neglect of space. In M. Posner and O. Marin (eds), *Attention and Performance*, vol. 11. Hillsdale, New Jersey: Erlbaum.
- Bisiach, E., Geminiani, G., Berti, A. and Rusconi, M. L. 1990: Perceptual and premotor factors of unilateral neglect. *Neurology*, 40: 1278-1281.
- Bisiach, E. and Luzzatti, C. 1978: Unilateral neglect of representational space. *Cortex* 14: 129-133.
- Bisiach, E. and Vallar, G. 1988: Hemineglect in humans. In P. Boller and J. Grafman (eds), *Handbook of Neuropsychology*, vol. 1. Amsterdam: Elsevier.
- Boden, M. 1990: *The Philosophy of Artificial Intelligence*. Oxford: Oxford University Press.
- Brewer, B. 1992: Unilateral neglect and the objectivity of spatial representation. *Mind and Language*, 7: 222-239.
- Cassam, Q. 1987: Transcendental arguments, transcendental synthesis, and transcendental idealism. *Philosophical Quarterly*, 37: 355-378.
- Cassam, Q. 1989: Kant and reductionism. *Review of Metaphysics*, 43: 72-106.
- Cassam, Q. forthcoming: Transcendental self-consciousness. In P. K. Sen and R. Verma (eds), *The Philosophy of P. F. Strawson*.
- Chisholm, R. 1981: *The First Person*. Brighton: Harvester Press.
- Collingwood, R. G. 1946: *The Idea of History*. Oxford: Oxford University Press.
- Coltheart, M. 1980: Deep dyslexia: a right-hemisphere hypothesis. In M. Coltheart, K. Patterson and J. C. Marshall (eds), *Deep Dyslexia*, pp. 326-380. London: Routledge and Kegan Paul.
- Davidson, D. 1984: What metaphors mean. In *Inquiries into Truth and Interpretation*, 245-264. Oxford: Oxford University Press.
- Davies, M. 1986: Tacit knowledge, and the structure of thought and language. In C. Travis (ed.), *Meaning and Interpretation*, 127-158. Oxford: Blackwell.
- Davies, M. 1987: Tacit knowledge and semantic theory: Can a five per cent difference matter? *Mind*, 96: 441-462.

- Davies, M. 1989: Tacit knowledge and subdoxastic states. In A. George (ed.), *Reflections on Chomsky*, 131–152. Oxford: Blackwell.
- Dennett, D. 1984: Cognitive wheels: the frame problem of AI. In C. Hookway (ed.), *Minds, Machines and Evolution*, 129–151. Cambridge: Cambridge University Press.
- Dennett, D. C. 1991: *Consciousness Explained*. Boston: Little, Brown.
- Diamond, S. 1972: *The Double Brain*. London: Churchill Livingstone.
- Evans, G. 1973: The causal theory of names. *Proceedings of the Aristotelian Society*, supp. vol. 47: 187–208.
- Evans, G. 1981: Semantic theory and tacit knowledge. In S. Holtzman and C. Leich (eds), *Wittgenstein: To Follow a Rule*, 118–137. London: Routledge and Kegan Paul. (Reprinted 1985 in *Collected Papers*, 322–342. Oxford: Oxford University Press.)
- Evans, G. 1982: *The Varieties of Reference*, ed. J. McDowell. Oxford: Oxford University Press.
- Fogelin, R. 1985: *Hume's Skepticism in the Treatise of Human Nature*. London: Routledge & Kegan Paul.
- Gallistel, C. R. 1980: *The Organization of Action: A New Synthesis*. Hillsdale, New Jersey: Erlbaum.
- Gallistel, C. R. 1990: *The Organization of Learning*. Cambridge, Mass.: MIT Press.
- Gazzaniga, M. 1988: In A. J. Marcel and E. Bisiach (eds), *Consciousness in Contemporary Science*, 226ff. Oxford: Oxford University Press.
- Gibson, J. J. 1979: *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Goldman, A. I. 1989: Interpretation psychologized. *Mind and Language*, 4: 161–185.
- Goldman, A. I. 1992: In defense of the simulation theory. *Mind and Language*, 7: 104–119.
- Goldman, A. I. 1993: The psychology of folk psychology. *Behavioral and Brain Sciences*, 16: 15–28.
- Gopnik, A. and Wellman, H. 1992: Why the child's theory of mind really is a theory. *Mind and Language*, 7: 145–171.
- Gordon, R. M. 1986: Folk psychology as simulation. *Mind and Language*, 1: 158–171.
- Gordon, R. M. 1992a: The simulation theory: objections and misconceptions. *Mind and Language*, 7: 11–34.
- Gordon, R. M. 1992b: Reply to Stich and Nichols. *Mind and Language* 7: 87–97.
- Gordon, R. M. 1992c: Reply to Perner and Howes. *Mind and Language*, 7: 98–103.
- Gordon, R. M. in press: Simulation without introspection or inference from me to you. In M. Davies and T. Stone (eds), *Mental Simulation: Philosophical and Psychological Essays*. Oxford: Blackwell.
- Harris, P. L. 1989: *Children and Emotion: The Development of Psychological Understanding*. Oxford: Blackwell.
- Harris, P. L. 1991a: The work of the imagination. In A. Whiten (ed.), *Natural Theories of Mind: The Evolution, Development and Simulation of Everyday Mindreading*, 283–304. Oxford: Blackwell.
- Harris, P. L. 1991b: Letter to Josef Perner, 30 May 1991.

- Harris, P. L. 1992: From simulation to folk psychology: the case for development. *Mind and Language*, 7: 120–144.
- Heal, J. 1986: Replication and functionalism. In J. Butterfield (ed.), *Language, Mind and Logic*, 135–150. Cambridge: Cambridge University Press.
- Heal, J. in press: How to think about thinking. In M. Davies and T. Stone (eds), *Mental Simulation: Philosophical and Psychological Essays*. Oxford: Blackwell.
- Hurley, S. L. in preparation: *The Reappearing Self*.
- Jeeves, M. A. 1965: Agenesis of the corpus callosum — physio-pathological and clinical aspects. *Proceedings of the Australian Association of Neurologists*, 3: 41–48.
- Johnson-Laird, P. N. 1983: *Mental Models*. Cambridge: Cambridge University Press.
- Kant, I. 1933: *The Critique of Pure Reason*. Tr. Kemp Smith, N. London: Macmillan.
- Lockwood, M. 1989: *Mind, Brain and the Quantum: The Compound 'I'*. Oxford: Blackwell.
- Marcel, A. J. 1993: Slippage in the unity of consciousness. In Ciba Foundation Symposium No. 174, *Experimental and Theoretical Studies of Consciousness*. Chichester: John Wiley.
- Marks, C. E. 1981: *Commissurotomy, Consciousness and the Unity of Mind*. Cambridge, Mass.: MIT Press.
- Milner, A. D. and Jeeves, M. A. 1979: A review of behavioural studies of agenesis of the corpus callosum. In I. S. Russell, M. W. Van Hof and G. Berlucchi (eds), *Structure and Function of Cerebral Commissures* 428–483. London: Macmillan.
- Nagel, T. 1979: Brain bisection and the unity of consciousness. reprinted in T. Nagel, *Mortal Questions*. Cambridge: Cambridge University Press. (First published in 1971 in *Synthese*, 20.)
- O'Keefe, J. 1985: Is consciousness the gateway to the hippocampal cognitive map? A speculative essay on the neural basis of mind. In D. A. Oakley (ed.), *Brain and Mind*, 59–98. London: Methuen.
- O'Keefe, J. 1990: A computational theory of the hippocampal cognitive map. In J. Storm-Mathisen, J. Zimmer and O. P. Ottersen (eds), *Progress in Brain Research*, 83: 301–312. Amsterdam: Elsevier.
- O'Keefe, J. 1991: The hippocampal cognitive map and navigational strategies. In J. Paillard (ed.), *Brain and Space*, 273–295. Oxford: Oxford University Press.
- O'Keefe, J. 1993: Kant and the sea-horse. In N. Eilan, B. Brewer and R. McCarthy (eds), *Spatial Representation: Problems in Philosophy and Psychology*, 43–64. Oxford: Blackwell.
- O'Keefe, J. and Nadel, L. 1978: *The Hippocampus as a Cognitive Map*. Oxford: Oxford University Press.
- Parfit, D. 1984: *Reasons and Persons*. Oxford: Oxford University Press.
- Peacocke, C. 1986: Explanation in computational psychology: language, perception and level 1.5. *Mind and Language*, 1: 101–123.
- Peacocke, C. 1989: When is a grammar psychologically real? In A. George (ed.), *Reflections on Chomsky*, 111–130. Oxford: Blackwell.
- Peacocke, C. 1992: *A Study of Concepts*. Cambridge, Mass.: MIT Press.
- Peacocke, C. 1993: Externalist explanation. *Proceedings of the Aristotelian Society*, 93: 203–230.

- Perner, J. 1991: *Understanding the Representational Mind*. Cambridge, Mass.: MIT Press.
- Perner, J. and Howes, D. 1992: 'He thinks he knows': and more developmental evidence against the simulation (role taking) theory. *Mind and Language*, 7: 72–86.
- Piaget, J. and Inhelder, B. 1951/1975: *The Origin of the Idea of Chance in Children*. New York: Norton.
- Powell, C. T. 1990: *Kant's Theory of Self-Consciousness*. Oxford: Oxford University Press.
- Quine, W. V. O. 1960: *Word and Object*. Cambridge, Mass.: MIT Press.
- Rorty, R. 1970: Strawson's objectivity argument. *The Review of Metaphysics*, 24: 207–244.
- Schwyzler, H. 1990: *The Unity of Understanding*. Oxford: Oxford University Press.
- Sergent, J. 1990: Furtive incursions into bicameral minds. *Brain*, 113: 537–568.
- Seymour, S., Reuter-Lorenz, P. and Gazzaniga, M. 1994: The disconnection syndrome: basic findings reaffirmed. Abstracted in *The Society of Neuroscience*, 1993.
- Shebilske, W. L. 1984: Context effects and efferent factors in perception and cognition. In W. Prinz and A. F. Sanders (eds), *Cognition and Motor Processes*. Berlin: Springer-Verlag.
- Shoemaker, S. 1984: Causality and properties. In S. Shoemaker, *Identity, Cause and Mind*. Cambridge: Cambridge University Press.
- Sperry, R. W. 1990: Forebrain commissurotomy and conscious awareness. In C. Trevarthen (ed.), *Brain Circuits and Functions of the Mind*. Cambridge: Cambridge University Press.
- Stich, S. and Nichols, S. 1992: Folk psychology: simulation or tacit theory? *Mind and Language*, 7: 35–71.
- Stich, S. and Nichols, S. in press: Second thoughts on simulation. In M. Davies and T. Stone (eds), *Mental Simulation: Philosophical and Psychological Essays*. Oxford: Blackwell.
- Strawson, P. F. 1959: *Individuals*. London: Methuen.
- Strawson, P. F. 1966: *The Bounds of Sense*. London: Methuen.
- Tegnèr, R. and Levander, M. 1991: Through a looking glass. *Brain*, 114: 1943–1951.
- Trevarthen, C. 1974: Analysis of cerebral activities that generate and regulate consciousness in commissurotomy patients. In S. Dimond and J. G. Beaumont (eds), *Hemisphere Function in the Human Brain*. London: Elek Science.
- Trevarthen, C. 1984: Biodynamic structures. In W. Prinz and A. F. Sanders (eds), *Cognition and Motor Processes*. Berlin: Springer-Verlag.
- Walker, R. 1978: *Kant*. London: Routledge.
- Wiggins, D. 1980: What would be a substantial theory of truth? In Z. van Straaten (ed.), *Philosophical Subjects: Essays Presented to P. F. Strawson*, 189–221. Oxford: Oxford University Press.
- Wilkie, D. M. and Palfrey, R. 1987: A computer simulation model of rats' place navigation in the Morris water maze. *Behavioural Research Methods, Instruments and Computers*, 19: 400–403.
- Williams, B. 1978: *Descartes: The Project of Pure Inquiry*. Harmondsworth: Penguin.
- Wilson, M. D. 1987: *Descartes*. London: Routledge & Kegan Paul.

- Wimmer, H. and Perner, J. 1983: Beliefs about beliefs: representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 13: 103-128.
- Wimmer, H., Hogrefe, G.-J. and Perner, J. 1988: Children's understanding of informational access as a source of knowledge. *Child Development*, 59: 386-396.