Proceedings of the British Academy, 83, 3–20

PHILOSOPHICAL LECTURE

Objects and Objectivity

JOHN CAMPBELL
New College, Oxford, OX1 3BN

1. The causal structure of physical objects

IT IS WIDELY HELD that the capacity for spatial thought depends upon the ability to refer to physical things. The argument is that the identification of places depends upon the identification of things; places in themselves are all very much alike and can be distinguished only by their spatial relations to things. So one could not so much as think about places unless one could think about things (Strawson, 1959). It has to be acknowledged that our identifications of places are greatly enriched by our ability to refer to physical things. But, as we shall see, it is possible to identify places without identifying objects. This raises the question whether there is any fundamental role that physical objects do play in our spatial thinking. I begin with the ways in which reference to physical objects enriches our capacity to identify places. We shall then consider whether reference to places as such demands reference to objects, and if not, what special role there might be for physical things in spatial thinking.

A physical object has a certain causal structure. We can bring this out by reflecting on the way in which the properties of a physical thing affect its behaviour. Some of the properties of a thing just are propensities for it to behave in particular ways in particular circumstances. For example, being elastic, or brittle, are dispositional characteristics, they say that the thing will behave one way rather than another under pressure. But other properties of a thing, such as its size and

Read at the Academy 13 March 1993.

shape, or the material of which it is made, are not dispositions to affect its behaviour in particular circumstances. These properties are not entirely irrelevant to the way a thing behaves, though — it is just that they are not directly tied to some one pattern of behaviour. Rather, the shape of a thing affects its propensities to behave, in a way that depends upon which other properties the thing has. So if something is bowl-shaped, it will hold liquids if it is made of glass, but not if it is made of wire mesh. A property of a thing affects its dispositions, conditionally upon its possession of other properties (Shoemaker, 1984). The dispositions of a thing to behave in this or that way in various circumstances are the upshot of its possession of a complex of properties. This network of properties of the thing, and the way in which they interact to yield the behaviour of the thing, constitute its having a certain causal structure.

This point about causal structure tells us something about how physical objects can function as common causes of correlated phenomena. And that makes it possible for us to know informative identities which are used in identifying places. Let us take this slowly. It can be of value to postulate a single object as the common cause of a number of correlated effects. For we may have a thing with a single striking property which combines with various permutations of its other properties to yield the various correlated effects. For example, this is how it might go in the case of postulation that various deaths are the work of a serial killer. We have a single striking characteristic, the psychosis, revealed by the string of deaths, which combines with various permutations of the other properties of the killer, such as his height, build, strength and dexterity to yield the various correlated effects, which we use in 'building up a picture' of their common cause. Now the capacity of objects to function as common causes characteristically underlies our ability to frame informative identities about the landmarks we use in finding our way around. Suppose I have recently arrived in a town, and have got to know a selected set of routes around a particular area. Then one day having gone a few yards off my usual path I am suddenly quite lost. What will orient me is my realizing that I am looking at an unfamiliar view of a familiar object — that the building in front of me is the Post Office, seen from a new angle. What is going on here is that the Post Office is functioning as the cause of both my earlier observations of it and my current perception. I need not go through this 'common cause' reasoning explicitly myself in grasping the informative identity. I do not have to reflect explicitly on

the causes of my observations. What grounds my acceptance of the identity is my registering a correlation between the properties of the Post Office as first seen, and the building I now observe. But not just any correlation of properties will do; not just any correlation of properties will be enough to give me knowledge of the identity. What condition is there on the correlation of properties which I must register if I am to have knowledge of the identity? The correlation will be sufficiently strong to ground knowledge of the identity if it is sufficiently strong to ground 'common cause' reasoning which would establish a single object as the cause of both sets of observations. And the identity judgement will be true only if the upshot of that 'common cause' reasoning would be correct — only if there is in fact a common cause of the past and present observations. So the capacity of objects to function as common causes is being exploited here, even though the subject does not explicitly engage in 'common cause' reasoning.

There are also informative identities knowledge of which does not in this way depend upon 'common cause' reasoning. These are identities which we come to know by tracing the spatio-temporal paths of objects of the same type, and finding that they are continuous. For example, the discovery that the Scarlet Pimpernel is Sir Percy Blakeney might be of this type, if one finds out by following Sir Percy and observing the transfiguration. And the discovery that the Morning Star is the Evening Star could be of this sort, if we do it by finding that the trajectories of the planets are continuous with one another. But this kind of procedure could not provide the dramatic aid to place-identification exemplified by the knowledge, 'that building is the Post Office', precisely because it depends on finding where things are at various times. It could not itself be used to generate a set of place-identifications on the basis of the informative identity; rather the place-identifications are grounding the identity. To enrich place-identification, we need knowledge of informative identities which does not itself depend upon knowledge of the places of the object in question at various times; and it is precisely the causal structure of physical objects that enables us to know informative identities in this way.

The fact that the propensities of an object to behave in one way or another depend upon an interrelated complex of its properties also tells us something about how the object is internally causally connected over time, and that in turn is what makes it possible for us to recognize objects as we do. The way an object is later causally depends, in part, upon the way it was earlier. Suppose we consider one object interacting

with another, in a collision, for example. The way the object is after the collision depends in part upon the way it was before — that is why it is worth designing cars with an eye to safety, so that they will crumple easily upon impact. But exactly how the object is after the collision will characteristically depend upon a complex of its properties, such as the shape, material and rigidity of its bodywork. This means that recognition of a thing need not be a matter simply of registering the presence of a number of characteristics which it also had last time one encountered it. It has to do, rather, with grasping how the remembered thing could after a number of collisions have become the currently perceived thing. This kind of imaginative recognition is perhaps exercized most strikingly in recognizing the now ravaged features of an acquaintance one has not seen these past ten years, grasping how the past person could have become this one after a number of collisions. But it can also be used in connection with landmarks, as when, returning to a town one has not visited for years, one recognizes the old school, now dilapidated and partly demolished, and realizes that all the old places are close at hand. So once again, the causal structure of physical objects means that the ability to refer to them can enrich our capacity for the identification of places.

To sum up. We have been looking at ways in which the capacity for place-identification is enriched by the ability to refer to physical things. The enrichment is owed to the causal structure of physical things. One dimension of this causal structure is the capacity of a physical thing to function as a common cause of various phenomena, and we put our grasp of this 'common cause' aspect of causal structure to work in our judgements of informative identity. These judgements of informative identity can serve as the basis for whole ranges of place-identifications. For example, on realizing that the mountain Afla is the mountain Ateb, one now knows that the places one identified by their spatial relations to Afla stand in those very relations to Ateb, and if one is now facing Ateb one can perceptually identify those places. But to do this epistemic work, this knowledge of an informative identity involving landmarks must not depend on knowledge of the spatio-temporal continuity of the landmarks, for that already requires knowledge of place. It must rather exploit the causal structure of the landmark, and the fact that it can function as a common cause of various observations of it. The other dimension of the causal structure of physical objects is their 'internal causal connectedness', which we put to work in the imaginative exercise of our recognitional abilities. Per-

ceptual recognition of an object is not always a matter simply of remarking properties shared by the past object and the present one; it may rather involve the ability to see the old object in the new, to see how the world could have, by acting on the old object, produced the new. This means that a landmark may change significantly without changing out of all recognition, so that places may continue to be identified by their spatial relations to it. And when we look in detail at the ways in which we exploit these two dimensions of causal structure in physical things, we see that they rest on the fact that the properties of an object do not in general affect its behaviour one-by-one, but rather in combination with one another.

These aspects of causal structure are not peculiar to physical objects; they are more abstract than that. A light ray, or a wave out at sea, can function as a common cause of various phenomena, and is internally causally connected over time. What is distinctive of ordinary physical objects is the particular complex of interrelated properties that they have. Centrally, these include shape, size, solidity, and motion or rest. But there is no concise definitive list to be given here: there are many cases in which we might well be unsure whether we have a causal structure which puts to work enough of the right family of properties to constitute a physical thing — we might, for example, consider cities, rabbit-warrens or waterfalls to be difficult cases. In this way the notion of a 'physical object' is the paradigm of a family resemblance term.

Although our use of physical objects really does enrich our ability to identify places, we shall see that it is not essential to it. This raises the question whether there is any aspect of our ordinary spatial thought that really does exploit our capacity to refer to physical things, conceived of as having this type of causal structure. We shall see that the temporal dimension of ordinary spatial thinking does put this causal structure to work. But let us take these points in order.

2. Place-identification without physical objects

As I said, I think it cannot be maintained that place-identification in general depends upon the reidentification of things. Place-identification is a much more primitive capacity. Consider, for example, a creature which has the ability to keep track of places by keeping track of its own movements. At any one moment it can find the vector from it to

a particular place, and it can update its grasp of what that vector is by keeping track of the distances and directions for which it moves itself. Or again, we might consider a creature which engages in dead reckoning using some external compass, such as the sun and the time of day. The animal uses the compass to keep track of its own various swoops and sallies, and can then use path integration to find the direct route home. An animal which keeps track of places by using its own movements in this way is certainly reidentifying places, but need not be thinking of the features by which it sets its compass as physical things.

These points might be acknowledged, and yet it still be held that when an animal uses landmarks to reidentify places, those landmarks must be physical things. But here we have to remark that there seems to be a level of thought more primitive than the level at which we have thought about physical objects. This primitive level of thought is perhaps exemplified by the way in which we ordinarily think about the stars. If we are asked, as we look at the night sky and try to identify constellations, whether we think of the stars as physical objects or as points of light with no more causal significance than shadows — tears in the fabric of the sky — then there may be no immediate answer. We are not really thinking of them as either; the question had not come up before. We were at a more primitive level of thought than that.

Let me expand on the distinction that is not being drawn here. The pool of light thrown by a projector onto a wall is not causally structured in the way that a physical object is. It is not internally causally connected over time, the way the pool of light is at one time does not have its earlier condition as a causal determinant, its condition is determined always by the contemporaneous state of the projector and the surroundings. And while a pool of light can function as a common cause, for example of various observations by different subjects, the properties with respect to which one can engage in common-cause reasoning in the case of a pool of light are severely restricted. This is because the pool of light is not the bearer of a complex of interrelated properties which jointly determine its behaviour, so that we cannot have the structure in common-cause reasoning that we earlier saw to be available in the case of a physical object, such as a serial killer. But for an organism navigating around its environment, and using landmarks to identify places, it does not matter whether the landmarks are physical objects or not. A stably located pool of light would do just as well. Or it could operate at the more primitive level of thought at which the

distinction between a pool of light and a causally structured physical object is simply not drawn.

Consider a very simple representative of a whole class of navigational systems. This is the 'triangulation' model used to explain the behaviour of rats in a water maze (Wilkie and Palfrey, 1987). Here the rats are placed in a swimming pool filled with an opaque liquid. There is a submerged platform to which they learn to make their way. The platform, being submerged in an opaque liquid, cannot be seen by the rat. But it can reliably make its way to it, from any starting-point in the pool, so long as it keeps its relation to the distinctive landmarks it can see around the pool. The 'triangulation' model supposes that what happens is this. Once on the platform, the animal records the distances to each of the cues it can see. Then when it next tries to get to the platform, it notes the distances from where it is to each of the landmarks around it. If the distance to a landmark is currently greater than it was from the goal platform, the animal swims towards it. If the distance is less than it was from the goal, the animal swims away from it. Its movement is the resultant of all these calculations. How must the animal be thinking of the cues hung around the pool? Must it be thinking of them as physical objects, or might it be thinking of them as more like shadows or points of light? There is no reason why the animal should have had to make up its mind about that. They are recognizable and stably at those places, and that is really all it needs.

There is some precedent for acknowledging a 'feature-placing' level of thought, more primitive than the level at which we have reference to things, though there has often been some uncertainty about whether features can be assigned locations at this level of thought. But there is no reason why not; our navigating animals can be taken to be registering the presence of 'red square at (d, α)', and so on, where the distance and angle are given in terms relating to its capacity for perception and action in the space. That the particular located feature is not being thought of as a physical thing shows up in the fact that the animal may be quite incapable of grasping an informative identity involving that located feature; and it may be quite incapable of recognizing it as the same thing over a period in which there is some change in its characteristics; it may be wholly incapable of working with the idea of the landmark as an object whose characteristics change when it is acted on. These located features are being thought of as causally inert; there is no answer being given to the question whether they are to be thought

of as shadows, or pools of light, or if they are rather causally structured processes, such as physical objects.

I began with an argument for the dependence of the place-identification on thing-identification which runs as follows. Places in themselves are all very much alike, so the only way in which they can be distinguished is by their spatial relations to physical things. In effect we have been considering two lines of objection to this argument. One is that places may be distinguished by their spatial relations not to objects, but to more primitive featural aspects of the environment. The other is that places may be distinguished by their various relations to the animal — where it has to look to attend to them, how it has to move to act on them. This does not require the animal to be self-conscious — it may exploit the fact of these relations to itself, rather than have thought about these relations to itself, in distinguishing places. An animal which keeps track of places by keeping track of its own movements need not be thinking of itself as a physical object — it need not be thinking of itself at all. The 'ability to keep track of its own movements' may simply consist in its ability to update the vectors to the targets around it. The animal operating in this way will be relying on spatial perception, as well as on its knowledge of its own movements — it needs spatial perception to find the vectors to places from which it begins. But it can have spatial perception without having perception of causally structured physical objects, rather than more primitive featural aspects of its surroundings.

There is a contrast between the cues which the animal uses as landmarks, in navigating around, and its targets in navigation. There is no need for the animal itself to interact with any of the landmarks, except by perceiving them. But the target, its destination through the navigation, typically will be something with which the animal interacts. It might be food, or a nest, or a mate, or prey, or its young. Here it does not seem right to say that the animal might as well be thinking of these things as shadows. It expects its young, once fed, to stay fed for a while; and when it eats, it expects this to produce some persisting effect on it. This does not reinstate the argument with which I began, because that argument was concerned with the need for places to be differentiated by their spatial relations to landmarks, and the present line of thought concedes that these landmarks need not be physical things. The point is rather that for the organism to have any use for spatial information it must have some causal expectations about its targets; it cannot be thinking of them as causally inert features. This is

not yet enough for the targets to be thought of as ordinary physical objects, however. For we do not yet think of them as having their behaviour dependent upon an interrelated complex of properties; there is as yet no grasp of the target as the bearer of a complex of properties any one of which affects its causal powers only conditionally upon its possession of the other properties. The organism need not be capable of grasping informative identities involving its targets — it thinks of the target as the bearer only of a single-track property, such as edibility, relating to its own interaction with it. And there may be no possibility of recognizing the target as the same through any change in the single-track property, a change occasioned by interaction between the target and the world.

We can, of course, also make sense of grasp of this richer notion of the targets as physical objects, which would show up in possession of a complex behavioural response to the objects, in that one's response to detection of any one of its properties depends upon which other properties one knows it to have. We can make sense of a spectrum of cases in which animals become increasingly alive to the possibility of informative identities involving their targets, and the possibility of recognition of them as the same again through more or less dramatic change in their characteristics. But this complex behaviour evidently is not a precondition of engaging in spatial thought as such. The identification of places does not depend upon the identification of things; the ability to identify places is more primitive.

3. Objective space

On the face of it, the position we have reached is puzzling. We think of spatial magnitudes as having a role to play in the physics of our environment. In particular, space is crucial as the condition of interaction: there is no action at a distance. For there to be interaction there must be contact. Now the interests of an animal using a navigational system are severely practical. It is not interested in geometry for its own sake. It is interested in it only for its physical implications. But the features in terms of which it is thinking are causally inert — it has not bothered even to make the elementary distinction between causally structured physical objects, and pools of light or shadows. How then can it be thinking of spatial properties and relations as causally significant? It is not in a position to think of contact as the condition

of interaction between two physical objects. And if we try to think of contact as the condition of interaction between two located features, we run up against the problem that we have no conception of interaction between features. What happens to something after an interaction is the joint upshot of the way it was before the interaction and the nature of the impact upon it. If we lose our grip on the idea that what happens to it later is causally affected by the way it was earlier — as we do at the primitive level of thought — then there really is no saying what an interaction would involve.

The problem, then, is to understand how an animal operating at the primitive level of thought might manage to assign causal significance to spatial properties and relations; in particular, how it might exercise grasp of the idea that contact is the condition of interaction. Of course, so far I have been stressing the point that the animal thinks of the features it uses as landmarks as causally inert. And as we saw, the animal does not think of its targets in that way. But the sense in which the animal does not think of the targets as inert needs to be specified with some care. It is not that the animal has to be able to think of its targets as capable of interaction with one another — it might perfectly well be representing places without having got so far. And, though here the point is subtler, it is not that the animal has to be thinking of its targets as things with which it interacts itself. It need not have any reflective understanding of the relations between itself and its surroundings. It need not be self-conscious at all, even though it is identifying places. Its grasp of the causally significant characteristics of its targets has to do rather with the way in which it acts upon them. What matters here is not the animal's thinking about the way in which it acts upon its surroundings, but simply the way in which it does intentionally act upon them.

When we reflect on this point, we begin to see the sense in which the animal can be said to appreciate that contact is the condition of interaction. The animal's grasp of that idea consists in the fact that it will attempt to get to something with which it wants to interact, and try to get away from something it does not want to interact with. Consider, for example, an animal using the 'triangulation' system to navigate in the water maze. It thinks of the landmarks hung around the pool as causally inert — they might as well be pools of light or shadows. But it does grasp the causally significant aspect of the target platform. What then does it come to, that the animal regards contact as the condition of interaction? Its grasp of the idea consists in the

fact that it will use the system to get to the platform, and it regards contact with the platform as the condition of using it.

This explains how the animal is giving physical significance to spatial properties and relations, even though it is representing its environment in terms of features rather than physical things. It gives physical significance to spatial properties and relations not by having an explicit reflective overview of the causal relations which hold amongst the various happenings in its environment, but by using its grasp of those spatial properties and relations in practice, to move through its space and to act in that world. This does not demand that the animal should be thinking in terms of causally structured physical objects, because the animal's grasp of causation is constituted not by a reflective grasp of its relations to its surroundings, but by the fact of its relations to its surroundings, its capacity for intentional action in its world.

This immediately suggests the possibility of a different way of thinking about space, one which really would demand the ability to refer to physical things. This is a way of thinking about space which does not depend upon one's own engagement in the space, one's own ability to intervene in it. One would acknowledge a physical significance for spatial magnitudes — in particular, contact could be acknowledged as the condition of interaction. But the interaction now need not be interaction between oneself and the surroundings, and even in so far as it is interaction between oneself and one's surroundings, it is the thought of such interaction, rather than actual engagement in it, that one would be using to give physical significance to the spatial magnitudes. One would have a 'reflective' or 'detached' mode of thought about one's surroundings. And thinking in this way, would one not really need to be thinking about physical objects? This kind of disengaged understanding of interactions would seem to demand a grasp of the causal structure of the objects which were interacting with one another.

On the other hand, the very idea of this 'objective' mode of spatial thought invites scepticism. There is a tradition of empiricist-pragmatist criticism of the notion, demanding that thought about places must always be understood in terms of one's own potential for action in the space. This tradition was vigorously resisted by John O'Keefe, most recently in his elaboration of the details of a genuinely allocentric or 'objective' type of spatial thought (O'Keefe, 1990, 1991; O'Keefe and Nadel, 1978). On this model, there are two preliminary stages in an animal's construction of a map of its environment. The animal has first

to identify the 'centroid' of its surroundings. This is a 'notional' point, in that there may be no distinctive physical feature at that place — it is the geometric 'centre of mass' of the environment. It is a fixed point, which does not move when the animal moves. Secondly, the animal finds a way of giving directions — it specifies the 'slope' of the environment, which is a gradient defined by the way that cues are distributed in the space. This direction stays static no matter how the animal turns, so the animal can specify which way it is going by specifying its angle with the slope. Suppose now that the animal records the vector from the centroid to each of its targets, using the slope to specify direction. This gives it a 'map' of its surroundings, which can be carried around and be of use wherever the animal is in its environment. If the animal ever wants to get to a specific target, what it has to do is to find the current vector from it to the centroid. Since it has the vector from the centroid to its target, it can now compute the direct vector from it to the target. There is evidently a considerable contrast between this model and the more primitive 'triangulation' system considered earlier. It is more powerful in which sets of geometric relationships between places in the environment it can represent, and it is not tied to a single target. But for present purposes, what matters is the similarity between the two models. On this model, there is still no need for the animal to be thinking in terms of causally structured physical objects. And the reason is, evidently, that the physical significance of the whole system is accorded to it by the animal through its use of the system in navigating through the space. We do not have here a detached or 'objective' way of understanding the physics of the space.

Most recently O'Keefe (1993) has acknowledged that the slope-centroid model does not qualify as a truly allocentric mapping system, and made some proposals about how it might function as a platform for the construction of such a system. Objects could be used by the system. We could allow for the ability to consider vectors from the animal to the centroid which do not represent its current position. We could allow for the progressive use by the animal of the distinction between agents and non-agents — the latter being preferable as landmarks — and grasp of its own properties as an agent. And we could allow for development by the animal of a richer physics of its surroundings and mental models to enable it to predict the actions of others. I have considerable sympathy with the direction of these proposals, but it is not entirely clear to me whether the animal is still thought of as having its grasp of the causal significance of this rich set of classifi-

cations exhausted by its own capacity for action and engagement in the space. If so, then we do not seem yet to have 'objective' spatial thought. If not, then we need to know at what point the frontier was crossed. But now it may seem that the empiricist-pragmatist critique is entirely in order. Is there any frontier here? Do we not rather have a set of types of representation, all located at a greater or lesser distance from action and perception, but all ultimately given meaning by their significance for one's engagements in the space?

The nerve of the notion of objectivity is the idea of a way of thinking about space that is detached, or disengaged, from the demands of perception and action: this way of thinking about space is not given significance by its implications for current and future perception and action. One way of thinking about objectivity is as a matter of achieving a more or less synoptic view of an area. We think of an objective view of the area as like an aerial photograph of it, and greater objectivity is achieved by progressively moving the photographer further and further away. This imagery plays a role in making it seem that objectivity is something towards which we can approximate but cannot finally attain, for there is no limit to the distance from which the photograph can be taken. But this is not the only way of thinking about objectivity. I want to propose that the best model is provided not by the aerial photograph but by the diary – a narrative of the events in a spatial region which may have no immediate implications for current perception and action, but is understood as an end in itself.

4. Representation of time

To understand the sense in which we ordinarily do achieve a 'disengaged' or 'objective' representation of the space we are in, we have to consider the temporal dimension of spatial thinking. And when we do this, we also see why causally structured physical objects have such a fundamental place in our thought. At first this second claim may seem surprising. There is considerable complexity and sophistication in the timing systems that animals may use, even though they show no ability to refer to physical objects (Gallistel, 1990). For instance, an animal may use an interval timer — an internal stopwatch — in controlling the time between its visits to a renewable food source. Too much time between trips means that someone else may get to the food source first. Too little time between trips means that the source may not have

been replenished. And there may be some complexity in how the animal uses its stopwatch to calculate which interval gives it the best rate of return. Obviously none of this demands an ability to refer to physical things, any more than does use of the navigational systems we considered earlier. So far, though, we do have only the use of an internal stopwatch. We do not yet have the ability to identify particular times. We do not yet have the use of a clock. So suppose we consider an animal which has a circadian clock. The clock may be an oscillator, something which simply repeats the same process over and over again, and the animal is sensitive to which phase of the cycle the oscillation is currently at. The oscillator may be kept in step with the external light-dark cycle. An animal using such a clock may develop a wide range of expectations as to what will happen at what phase of the day. And it still need not be thinking in terms of causally structured physical things. Although it keeps its oscillator in step with the external light-dark cycle, it need not represent the periods of light and dark as anything more than causally inert features of its surroundings.

The crucial point, though, is that we do not yet have representation of particular times. All we have is the animal's orientation with respect to the phase of the oscillation. It forms hypotheses about what happens at a particular phase of the oscillation, confirms, disconfirms or acts upon these hypotheses, and that is all. It has no need for, or use for, the idea that there are temporal relations between different events which occur at the same phase of the oscillation. Even if the animal uses a decay or accumulation process to distinguish between different events which occur at the same phase of the cycle, that does not mean it is regarding them as temporally related to each other.

The point is that what gives us the right to regard the animal as using the oscillator to record the times of various events is the use that the animal makes of the information on future occasions — when that phase of the cycle comes around again. This means the animal cannot be ascribed representation of anything more than phase. It could do nothing with the further information that the event happened at a particular past time, one among many particular past times all temporally related to each other, and which may or may not be occurring at the same phase of the cycle. Information about particular past times has no role in the animal's future engagements with its world. This is one of the basic differences between particular times and particular places, that there is no such thing as reidentifying a particular time, whereas there is such a thing as reidentifying a particular place.

Because of this, a grasp of time which is ultimately exhausted by its significance for the demands of action can be at most an ability to represent phase, since the same phase can be re-encountered and that can affect one's actions. It cannot be a representation of particular time.

What then is required for representation of particular times? We need to understand how that representation might be put to work in one's thinking; and for that, we have to consider the construction of a narrative of the events in one's environment which is 'disengaged' from the demands of action. In that sense, it is an 'objective' narrative. We have to consider the organism constructing a narrative of the events in its environment, where the causal relationships among these events are represented in a way that is internal to the narrative, and not exhausted by its implications for the present and future actions of the creature constructing the narrative. And it is in the construction of such a narrative that we see the fundamental role of physical objects in our thought.

In constructing a narrative of this kind, the times of events have to be given significance for the causal relationships between them. We need the notion of a process, whose later stages are causally dependent upon its earlier stages, so that the relative times of the various stages can be seen to matter for the causal relationships among them. In constructing a narrative of this kind, we evidently cannot rest with the representation of causally inert features — representation which is neutral on whether the items represented are more like shadows than things. Here we do have to consider causally significant particulars.

As I said earlier, there are two dimensions to the causal structure of a physical thing. There is the fact that it is internally causally connected over time, that the way it is later depends in part on the way it was earlier. And there is the capacity of the object to function as a common cause, entering into many interactions. This raises the question why the representation of particular times should have to deal with particulars which have just this kind of causal structure. Why would it not be enough simply to consider a number of causally related events, without bringing physical objects into the picture at all?

There may be many narratives to be constructed of the events in one's surroundings, many stories to be told of sequences of events which while having their own internal dynamics are more or less insulated from one another, as the story of one's professional life may be relatively if not completely insulated from one's progress as a

hockey player. Now one way in which we relate events happening at different times is by noting the phase of some cycle at which they occurred. Once we have the possibility of a process having its earlier stage at a phase of the cycle, and its later stage at the same phase of the cycle, we can use this causal connection to distinguish between different particular times occurring at the same phase, and then use the cycle itself in temporally relating various particular events. But we would have no use for this procedure if there were not a rich set of causal relations among the various events in the narrative. If the narrative were highly fragmented, there would be no point to the assignation of temporally related particular times to all the events in the narrative.

The causal structure of physical objects plays a crucial role in securing the unity of the narrative of events in one's surroundings. The two aspects of this causal structure are that objects function as common causes, entering into many interactions, and that they are internally causally connected over time. The fact that it can be one and the same object which enters into various interactions plays a crucial role in causally connecting those various events, and hence giving some significance to the temporal relations between them. But evidently sameness of object could not do this work unless objects themselves were internally causally connected; unless an object entering into one interaction bore upon it the marks of its earlier interactions. It is those two aspects of the causal structure of physical things that means they can play a fundamental role in making it possible for us to construct 'detached' or 'objective' narratives of the history of our environment.

5. Self-determination

I have been emphasizing that there is a 'disengaged' or 'objective' level at which one's grasp of the causal significance of these modes of spatio-temporal thought is not exhausted by its implications for one's future perceptions and actions. This raises the natural question: what then is the use of it? What is the point of engaging in thought at this level? Pursuit of this line of thought is one source of the kind of empiricist-pragmatist criticism of the notion of objective space which I mentioned earlier. To address it, I end with some remarks about the bearing of these points on the notion of a person.

Persons are concrete objects, and they are causal unities. What is distinctive about them is that we think of them as causally structured

at the level of their psychological properties. The psychological properties of a person may operate in the causation of a range of social phenomena; people function as common causes at the level of psychological properties. And people are internally causally connected at the level of their psychological properties, most distinctively in the way that we think of memories as marks left on us by earlier perceptual interactions with our surroundings. Now self-consciousness, knowledge of oneself as an object, requires that one have some grasp of one's own causal structure: that is part of what gives one the right to use the first person as a referring term. But what then is involved in grasping one's own causal structure at the level of psychological properties? What is required is that one should not just think in terms of one's psychological life as a sequence of states, as a shifting kaleidoscope of the mental, but that one should think of the states as the products in part of interactions between oneself and one's environment, and as themselves determinants in such interactions.

Space has a fundamental role in our explicit physics of our environment, because we think of spatial contact as the condition of causal interaction. Spatial thinking plays a similar role in understanding our perceptual interactions with the world. In the case of touch the condition is again spatial contact — to touch the thing you must touch it. But in general the condition of causal interaction in perception is not contact. The analogue of 'no action at a distance' is provided by our grasp of the enabling conditions of perception. For example, to see something one must be appropriately located with respect to it, looking in the right direction, and there must be nothing in the way. So in both the physical and the psychological cases, one's understanding of the conditions of causal interaction is provided by one's reflective understanding of spatial relations. This reflective understanding of spatial relations is not to be explained in terms of its direct implications for current perception and action, precisely because it relates to one's past interactions with one's surroundings. One thinks of one's memories as the traces left by, for example, past perceptions. So we must here be thinking in terms of an 'objective' space, distanced or 'disengaged' from the demands of current perception and action.

The question set by the empiricist-pragmatist critique is to explain what practical value it could have to possess such an 'objective' conception of space. But this poses the wrong question. We cannot simply take it for granted that humans have a particular range of goals, and that any mode of thought they may engage in can be seen as contribu-

ting to one or another of those goals. Humans do not just have a variety of wants and desires. One can have a particular conception of one's own life, what its shape is, how it has gone so far and how one wants the rest of it to go. This conception is the understanding of one's life as having a narrative, as being bounded by particular times, within which boundaries one interacts with others and one's surroundings. And that conception is always up for discussion and criticism and revision. It is in that determination of the guiding lights of one's own life that we find the most distinctive use of the kind of reflective thinking I have been describing.

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 UniversityPress.

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.

Schwyzer, 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.