



PAST,
SPACE,
&
SELF

JOHN
CAMPBELL



Past, Space, and Self

Representation and Mind
Hilary Putnam and Ned Block, editors

Representation and Reality

Hilary Putnam

Explaining Behavior: Reasons in a World of Causes

Fred Dretske

The Metaphysics of Meaning

Jerrold J. Katz

A Theory of Content and Other Essays

Jerry A. Fodor

The Realistic Spirit: Wittgenstein, Philosophy, and the Mind

Cora Diamond

The Unity of Self

Stephen L. White

The Imagery Debate

Michael Tye

A Study of Concepts

Christopher Peacocke

The Rediscovery of the Mind

John R. Searle

Past, Space, and Self

John Campbell

Past, Space, and Self

John Campbell

**A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England**



MIT Press

0262531313

Campbell
Past, Space, And Self



First MIT Press paperback edition, 1995
© 1994 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, and information storage and retrieval) without permission in writing from the publisher.

This book was set in Sabon by The Maple-Vail Book Manufacturing Group and was printed and bound in the United States of America.

First printing.

Library of Congress Cataloging-in-Publication Data

Campbell, John, 1956—

Past, space, and self / John Campbell.

p. cm. — (Representation and mind)

“A Bradford book.”

Includes bibliographical references and index.

ISBN 0-262-03215-5 (HB), 0-262-53131-3 (PB)

1. Space and time. 2. Mental representation. 3. Schemas (Psychology) 4. Egoism. I. Title. II. Series.

BF467.C35 1944

153.7—dc20

93-14038
CIP

For Sarah

Contents

Preface ix

Introduction 1

1 Frames of Reference 5

1.1 The Distinction 5

1.2 Egocentric Frames 8

1.3 Maps 19

1.4 Physical Objects and Objective Space 25

1.5 Common Causes and Informative Identities 33

2 The Past 37

2.1 Temporal Orientation 37

2.2 Primitive Physics: Causal Indexicality and Working
Concepts 41

2.3 Timing Systems 51

2.4 An Explicit Physics 59

2.5 Common Causes and the Time Order 64

3 The First Person 73

3.1 Trading on Identity 73

3.2 Keeping Track 82

3.3 The First Person as a Token-Reflexive 101

4 Self-Reference and Self-Knowledge	109
4.1 The Problem	109
4.2 The Perceptual-Demonstrative Model	121
4.3 Conceptual Role, Reference, and the Body Image	134
4.4 Psychological Structure	147
5 The Reductionist View of the Self	155
5.1 What Reductionism Is	155
5.2 Ascribing Experiences	160
5.3 Causal and Normative Relations	173
5.4 Fission and Death	184
6 Conceptual Structure and Self-Consciousness	201
6.1 A Repertoire of Concepts	201
6.2 Truth, Structure, and the Simple Theory	207
6.3 Realism	215
7 The Realism of Memory	223
7.1 The Opposition to Realism	223
7.2 Truth-Value Links	226
7.3 The Stepwise Conception of Memory	233
7.4 Incommensurability	240
7.5 Immersion	246
Notes	253
References	257
Index	267

Preface

This book began ten years ago with a conviction that an understanding of the first person has a foundational place in the theory of meaning and metaphysics. It has taken a long time to reach an articulation of the idea. In April 1990, I visited the project on Spatial Representation at King's College, Cambridge, and worked up the basis for this material. I owe an immense amount especially to Naomi Eilan for her sense of direction and uncompromising support, as well as to the other members of the project: Bill Brewer, Tony Marcel, Roz McCarthy, and Ian McLaren. Through the project, I met John O'Keefe, whose combination of zest and deep learning made it a pleasure to work on space.

What delayed me more than anything else was my acceptance of the conventional view that there is a dichotomy between an understanding of the first person as a piece of language governed by the rule that any token of 'I' refers to whoever produced it, on which it has nothing to do with self-consciousness, and an understanding of it as a kind of quasi-perceptual demonstrative meaning something like 'this person', on which it does express self-consciousness. Long discussions with Philippa Foot, beginning in 1988, were what forced me to abandon that view, and I benefited a great deal also from

discussions on the topic with Quassim Cassam and Timothy Williamson.

Over the years I have learned much from Charles Travis, whose strictures on context dependence I have tried to observe. Philosophers vary in how much they learn from those around them. I have always been aware of learning a great deal. In particular, I must thank Justin Broackes, David Charles, Jonathan Glover, Carsten Hansen, Susan Hurley, Gavin Lawrence, Michael Luntley, Penelope Mackie, Nick Rawlins, and Joanna Woodall, as well as the MIT Press and its reviewers. In 1991, I visited the Australian National University in Canberra and while there prepared the penultimate version of this material. The warmth of Philip Pettit and Eileen McNally made it a very pleasant trip. I completed this version in perfect surroundings, thanks to the generosity of Carsten and Georg Hansen.

I have kept my principal debt till last. Through discussion and through his books, Michael Dummett has affected my conception of what philosophy is and the depth and energy with which it is possible to pursue it.

Past, Space and Self

Introduction

In Somerset Maugham's *Of Human Bondage* the hero is given a densely patterned Persian carpet and assured that in its design lies the solution to the problem of meaning. Naturally, he examines the carpet closely, searching for encrypted messages or hidden pictures. Eventually it dawns on him that the point is not any individual detail but the whole design. One has to weave all the details of one's life into a meaningful pattern. This book is about one of the rudimentary prerequisites of achieving that sense of pattern: self-consciousness. Self-consciousness involves at least the capacity to think in terms of the story of one's past life. It is also a prerequisite of ordinary social interaction. Indeed, one of the most basic kinds of social interaction, from a developmental point of view, is the sharing of memories between parent and child. This kind of social bonding seems to be part of what explains the evolutionary benefit, the selection value, of autobiographical memory.

Self-consciousness, as I describe it, is not just a matter of thinking of oneself as temporally extended, so that one has a sequence of states one after the other. It involves grasping one's own causal structure. There are two dimensions to this grasp of causal structure. There is grasp of the idea that one's later states causally depend, in part, on one's earlier states: when

2 *Introduction*

one interacts with one's surroundings, what happens is the joint upshot of the character of the things around one and the way one is oneself. The other dimension in grasp of one's own causal structure is the idea that one can function as a common cause of various correlated events around one. Schematically, one topic that will concern us is the relation between knowledge of one's physical characteristics and knowledge of one's psychological characteristics in grasping one's causal structure.

The project of the book is to determine the relation between, on the one hand, grasp of the first person, which is the most immediate expression of self-consciousness, and, on the other, the ability to think about the spatial layout of the environment and one's place in it and the ability to think about time as extending forward and backward, the ability to orient oneself in a temporal frame of reference.

What is the plan of the book? I begin by describing the spatiotemporal framework presupposed by conceiving of oneself as causally structured in two dimensions. There is a kind of objectivity possible in spatial thinking that has to do with the possibility of assigning causal significance to spatial properties and relations at a reflective level removed from the immediate practical demands of interaction with one's surroundings. This kind of reflective or objective spatial thinking turns out to be demanded by thought about the past, about the history of oneself and one's surroundings. These first two chapters describe the basic spatiotemporal framework that we use. I then turn to the first person to try to determine how understanding it depends on grasp of that framework. I describe the sense of the first person and what it means to say that there is such a thing. I consider the relation of an understanding of the first person to grasp of the two dimensions of one's causal structure and also consider its connection with self-knowledge and ordinary social interaction. Some have held, contrary to the whole thrust of this

book, that the concept of the self is in fact not fundamental in our thinking and that it is better dispensed with altogether. So I look at the bases of this reductionist view. Then I draw things together and set out the way in which grasp of one's causal structure demands use of the spatiotemporal framework.

In describing these basic modes of thought about oneself and one's world, I am describing a cluster of conceptual abilities fundamental to ordinary thinking. Indeed, they define what conceptual thinking is: I argue that the ability to engage in these basic patterns of thought is what sets apart a conceptual thinker from an animal capable only of more primitive types of representation or simply of different types of representation. Possession of this core set of abilities defines possession of concepts. Throughout the book it would have been possible to use a systematic notation to differentiate human thinking from the representational methods used by other animals. I have not done this, partly because it is always obvious whether or not I am talking about human thinking, so there is no important ambiguity, and partly because one main aim of this book is to earn the right to that distinction, so it would have been doctrinaire simply to lay down from the outset how the distinction is to be applied. I hope the reader who already believes in a firm distinction between human conceptual thinking and other types of representation will not be unduly depressed by the fact that the distinction is not presupposed in what follows; ultimately, the point is to provide a foundation for the distinction.

The aim of this book is to show that a realist view of the past is presupposed in the exercise of these basic conceptual abilities, the conceptual skills constitutive of self-consciousness. This line of argument gives a different place to self-consciousness in philosophy from that traditionally given to it. The classical reason for giving a central place to self-consciousness was to respond

to skepticism: self-consciousness seemed to be a foundational point from which one might begin. The resulting notion of self-consciousness, restricted to aspects of the self about which one could not be wrong, bears little relation to self-consciousness as an everyday phenomenon and is of little interest outside that particular epistemological debate. I am proposing a different role for self-consciousness. It can be viewed as setting certain fundamental norms of thought. It is the basis of the demand for causal structure in one's conceptual organization, causal structure that is distinctive of possession of concepts. And it demands a realist view of the past.

The fact that a realist view of the past is demanded by self-consciousness does not of itself show that realism about the past is correct. This would involve a kind of background idealism: supposing that the form of the world is somehow an artifact of the demands of self-consciousness. But it has to be acknowledged that very general facts about the world can be discovered by other means. For example, perhaps quantum mechanics will ultimately be seen to have implications for our understanding of the past. Perhaps quantum mechanics will even show that the kind of reflective understanding of our place in the world that I describe in this book cannot in the end be achieved. But that would not show that we should simply set aside realism about the past and self-consciousness. Rather, a profound conflict would have been generated. Even though there is no guarantee that the demands of self-consciousness will be met by the world, whether to be self-conscious is in practice not a matter about which we have any choice. For as I began by remarking, self-consciousness is a precondition of ordinary social interaction. And it is needed if one is to find any kind of narrative meaning in life.

Frames of Reference

1.1 The Distinction

There is a distinction that philosophers and psychologists have tried to draw between different ways of thinking about space, about particular spatial regions. It is sometimes called, and I will call it, the distinction between *absolute* and *egocentric* space. But it is not a distinction between different types of regions. It is a difference between ways of representing, or thinking about, a particular region. The distinction is at best very indirectly related to the classifications physicists make of theories of space-time as relativistic or absolute. For that reason the word 'absolute' is unfortunate. But it is what the literature uses, and I hope the reader will be able to set these echoes aside.

Intuitively, the distinction is between thinking about space as a participant, as someone plunged into its center, as someone with things to do in that space, on the one hand, and, on the other hand, thinking about the space as a disengaged theorist. Any animal that has the relations between perception and behavior needed to direct action at particular places, to reach for things it can see, must be capable of this egocentric spatial thinking. But the more detached absolute conception is not so easily available. The distinction is between the way of thinking

of the space one is in that one uses when sitting at a dinner table, moving and acting in that space, and the way of thinking of the space used subsequently by the detective who tries to reconstruct the scene and to establish who did what. It is the distinction between thinking about the space from a particular point of view, as a subject at the center of one's world, and thinking about the space independently of any particular viewpoint on it, in an impersonal or absolute way.

The idea of absolute space sometimes appears in discussions of self-consciousness. When self-consciousness is conceived in this way, it can seem dizzying. What it demands is that one should build up a synoptic picture of the world, one that wholly abstracts from one's own place in the throng, and then somehow identify one of the people so pictured as oneself. What is dizzying is the kind of complete objectivity, the degree of abstraction from one's own busy concerns, that is required. A first interpretation is that what is wanted is a kind of top-down view, so that we think in terms of a kind of aerial photograph, and then one has to identify oneself as one of the people shown in the photo. But that would not be enough, for it would only give the viewpoint of the photographer, and we need a picture of the world that is objective, in that it is not from any viewpoint at all. It is a picture not drawn from any perspective. A natural reaction to this demand is to move the photographer further away from the earth, but of course no finite amount of movement will be enough. And once he is at an infinite distance, what will he be able to photograph? What would you expect to be able to discern from nowhere?

One reaction to the vertigo here would be to say that we do not need any kind of objective conception in order to be self-conscious. Immersed, as we are, in the thick of things, we have no need or use for such a conception in our everyday lives,

whereas self-consciousness is commonplace. So an objective conception is not demanded by self-consciousness. But this reaction leaves us in the dark about what self-consciousness might be. The real problem is not the notion of an objective conception as such but the way of interpreting the spatial imagery that leads to the dizziness just described. Self-consciousness does demand a kind of objectivity and does use a conception of absolute space. But to understand the notion of objectivity involved, what we have to look at is our ordinary thinking about time. What matters is the way we think about the historical order of events that have happened and are going to happen and about our own temporal location in that order—the span of our lives within it and where things stand now. We will see that this really does involve a kind of objectivity and does use a conception of absolute space, and it certainly matters in ordinary life.

It is part of the distinction between absolute and egocentric space that the absolute conception should not be somehow reducible to, or definable in terms of, the egocentric spatial thinking used in perception and action. For this reason, the distinction has characteristically been rejected by philosophers of an empiricist or pragmatist bent, who think that all ways of representing space must be explicable in terms of their connections with perception and action, in terms of their relations to egocentric space. In this chapter we will first look more critically at the question of how to characterise an egocentric frame of reference. We will then look at John O’Keefe’s specific description of an absolute frame of reference and determine how well it stands up to empiricist-pragmatist criticism. We will see that a notion of absolute space can ultimately be sustained and that it is used in our ordinary thinking about time. One question we will consider is whether, in constructing such an absolute or objective frame of reference,

physical objects play any role; they certainly do not seem to be needed in constructing an egocentric frame.

1.2 Egocentric Frames

Let us begin, then, with egocentric spatial thinking. The frame of reference that we use to identify places in directing our actions, in deciding where to move to, is an egocentric frame. On the face of it, an egocentric frame is a body-centered frame or one centered on a part of the body. The developmental psychologists Herbert Pick and Jeffrey Lockman (1981) put the idea as follows. They define a *frame of reference* to be “a locus or set of loci with respect to which spatial position is defined.” *Egocentric* frames of reference then are those that “define spatial positions in relation to loci on the body.” They are contrasted with *allocentric* frames of reference, “which simply means that the positions defining loci are external to the person in question” (1981, 40). This definition seems indeed to give a reductive account of the notion of an egocentric frame, defining it in terms of notions that genuinely seem to be more fundamental than it.

It is worth reflecting on the general form of the definition. In trying to say what is characteristic of an egocentric frame of reference, I am not dealing with a problem in pure mathematics. It is not, for example, on a level with the question of whether a frame of reference uses polar or Cartesian coordinates. In purely formal terms, the best we could do would be to say that it must be possible, using an egocentric frame, to specify spatial relations to a single privileged point, but that would not separate an egocentric frame from one centered on the sun, for example. We have to say something about the physical significance of the origin of the frame: we want to say,

for example, that it must be centered on the subject. This notion of the subject is not a purely formal notion of pure mathematics. But saying where the frame is centered is only one way of giving physical significance to the formal notions. An alternative would be to consider the physical significance of the axes of the frame of reference and to take them as fundamental. An egocentric frame would then be one whose axes had a particular kind of physical significance. It would then be a substantive thesis, rather than a definition, that egocentric frames are invariably centered on the body, or a part of the body. And it would be quite wrong, on that approach, to suppose that in using an egocentric frame, one must be identifying places by their relations to a body already identified.

The definition of an egocentric frame as a body-centered frame takes for granted the general notion of an object-centered frame of reference, and it says that the egocentric frames are a particular class of object-centered frame, namely, those centered on the body or a part of the body. The general notion of an object-centered frame is certainly legitimate. Consider an object such as a table or a bus. We can think of the internal spatial relations between its parts. We can use this system of internal spatial relations to identify points within the object. There may be natural axes that the object has. For example, given a pillar-box, we could define a set of axes by reference to its long axis, its line of sight as through the slot, and its coronal plane as through which the door moves when it is opened. So far what we have is a way of identifying points internal to the object. But the system of spatial relations that I have set up between the parts of the thing can be further used to identify points external to it. We could, for instance, identify a coconut on a palm tree as lying on a line through the bottom of one leg of a table and the top of another, and a hundred yards distant

in the direction going from bottom to top. This way of identifying places need not be used only with inanimate objects as its basis. One could equally well take the internal spatial relations between the parts of a horse, or its natural axes, and use them to give fully allocentric identifications of the places around it. One could do the same with a human body; one could do the same with what is in fact one's own body. And then, by Pick and Lockman's definition, what we have is an egocentric frame. Evidently, there is a finer distinction that we want to make here. For it is not as if we can assume extensionality: not just any way of thinking of the subject will do. The notion of egocentric space is a psychological notion; the reason we want it is to explain why the infant, for example, turns one way rather than another. In particular, perceptual knowledge of the body will not do. Merely seeing one's own body in a mirror, for example, and using it to set up a system of axes will not provide one with an egocentric frame.

The obvious proposal is that the subject has to be using direct, nonobservational knowledge of his own body constituted by his possession of a body image. In one use of the phrase 'the body image', it has to be thought of as referring to a relatively long-term picture of one's own physical dimensions. So someone's body image might be changed as a result of their having a skin graft or the loss of a limb or simply by growing up. In this use of the term, one's body image provides one with a general sense of what kinds of movement are possible for one. It assigns a particular structure to a creature that underlies its possibilities of movement. We cannot directly use this structure to set up a system of axes: it assigns no particular shape to the body. What we need is rather what O'Shaughnessy calls the "here and now" body image, which "is given by the description or drawing or model one would assemble in order to say how

the body seems to one *at a certain instant*. For example: torso straight, right cylindrical arm stretched out from body, crooked at right angles, etc.” (1980, 241). If a person has such a body image, he can plot the spatial relations between the various parts of his body and use them to construct a body-centered set of axes that will indeed be the egocentric axes. This proposal relies on a direct relation between the subject’s body image and his ability to act. We have to think of the body image as giving the subject a practical grasp of the ways in which it is possible to act, the possibilities of movement open to him. Of course, there must be some relation between these two conceptions of a body image. The immediate problem, though, is to understand why this shift, from outer perception of the body, such as seeing it in a mirror, to inner perception, as provided by a short-term body image, should be thought to achieve anything. After all, as we saw, simply managing to use the spatial relations between the parts of the body to set up a system of axes does not in general secure one an egocentric frame. Why should we think that an egocentric frame is guaranteed if one relies on the spatial relations between the parts of the body given in inner perception? The point here is that there is in general no direct connection between the mere use of an arbitrarily chosen body to set up coordinate axes and the subject’s capacity for directed spatial action. What the present proposal relies on is a direct relation between a subject’s short-term body image and his ability to act. We have to think of the short-term body image as giving the subject a practical grasp of the ways in which he can act, the possibilities of movement open to him. The reason this seems promising is that the short-term body image has direct connections with action of the type possessed by the egocentric axes. The proposal is that we can view the direct connection between action and egocentric space

as a product of the direct connection between action and the short-term body image. But now we have to ask how the body image has this direct connection with action. And we immediately face a dilemma. For how are the spatial relations between the parts of the body given in the body image? One possibility is that they are given in egocentric terms: one foot is represented as to the right of another, below the rest of the body, and so on. But then it can hardly be held that the subject uses the natural axes of his body to set up the egocentric axes; rather, he already has to use the egocentric frame to grasp the spatial relations between the parts of his body. Alternatively, suppose that the spatial relations between the parts of one's body are given in nonegocentric terms. Then there is no prospect of using the axes of one's body to set up an egocentric frame; one is in no better a position to do this with respect to the body of which one has inner perception than one would be with respect to a body of which one has outer perception. In both cases the problem is the same. One's grasp of egocentric spatial axes, with their immediate connections to moving and acting, cannot be generated from a grasp of spatial relations that are nonegocentrically given. Grasp of egocentric spatial axes must be taken as primitive.

This means that a certain kind of reductive ambition for the definition of an egocentric frame as a body-centered frame has to be abandoned. We cannot view this definition as explaining the notion of egocentricity in more fundamental terms. We cannot see it as defining egocentricity in terms of the generic notion of an object-centered frame of reference plus the notion of a body-centered frame. For when we inquire into the needed notion of body-centeredness, it turns out that it already appeals to the notion of the body as given in the body image, with its spatial relations given egocentrically. In particular, then, we

cannot take the body image to be more fundamental than the egocentric axes; we cannot derive them from it. The egocentric axes have to be taken as primitive, relative to the body image.

It might be asked whether the body image is not at any rate coordinate with the egocentric axes, so that they have to be taken as equally fundamental for spatial action. But while some egocentric reference frame is evidently essential if we are to have spatial action—otherwise the action could not be regarded as spatial at all—it does not seem that a subject needs to have a body image to be capable of egocentric spatial action, action we would want to explain by appealing to his possession of an egocentric frame of reference. Coordination and direction of spatial action may be achieved by purely distal specifications of the locations that are the endpoints of the actions, without the subject's having a single central body image at all (Scott Kelso 1982). If a body image is superimposed on the subject's egocentric axes, this is in addition to the requirements for him to be thinking about places egocentrically. So when the subject is identifying places egocentrically, he cannot be thought of as doing so by first identifying a physical thing, himself, through a body image and then identifying places by their relation to his body. Rather, his capacity to use the egocentric axes is more fundamental than his capacity to think in terms of a body image. The egocentric identification of places does not depend on a prior identification of a body. The notion of an egocentric frame is more fundamental than the relevant notion of body-centeredness. It is only when we have elucidated the notion of an egocentric frame that we are in a position to say what this notion of body-centeredness is.

How, then, are we to characterize egocentric frames of reference? One alternative approach would be to say that an egocentric frame is one defined by the axes up, down, left, right, in

front, and behind, with the origin identified as here. Places cannot be identified by directions from a single origin alone. We have to add something about the way distances are measured in this frame of reference, or at least we need some kind of order relation. Even so, this approach would not give us enough to say in general what an egocentric frame is. We want to allow as intelligible the hypothesis that humans may use many different egocentric frames. Consider, for example, the axes defining the movements of the hand in writing. There is no reason to suppose that this will be the very same set of axes used to define the movements of the whole body. Nevertheless, it is still an egocentric frame. So an approach that tries to define what it is for a frame to be egocentric by simply listing a particular set of axes will not work. Again, there is no reason to suppose that all species will use the same egocentric axes. For example, creatures that are jointed differently from us or that live deep underwater may use different axes. Finally, even if, by listing a suitable set of axes, we could give an extensionally correct identification of the egocentric frames, we would still have the explanatory work to do. We would still have to explain what it was about the terms 'left' and 'right', for example, that made them particularly connected to moving and acting, for example. Even so, the right way to give a general definition of the notion of an egocentric frame of reference may still be by defining a class of axes, rather than by making a general demand about where the frames must be centered. And, of course, we expect that an extensional approach here will not succeed: we have to grasp how the subject is apprehending those axes.

The axes distinctive of an egocentric frame are those that are immediately used by the subject in the direction of action. They may include, but need not be confined to, the natural axes of

the body. In the case of the in-front and behind axis we have a distinction defined in terms of the body and its modes of movement and perception. Its application to us depends on exploiting ways in which we are not symmetrical. If we were symmetrical, being double-jointed and able to look either way, then our current notions of in front and behind simply could not be applied to ourselves, could not guide our actions in the way that they do. But we are not symmetrical in this way, and the distinction does guide our actions. In the case of the up and down axis, we seem not to have here a distinction defined in terms of asymmetries of the body. It has to do rather with orientation in the gravitational field. The extensive apparatus we have to tell us how we are oriented in the gravitational field is precisely the apparatus that tells us which way is up. The reason this matters to us is, of course, the pervasive influence of gravity on every aspect of our ordinary actions. So here we have an egocentric axis not defined as a natural axis of the body. Of course there is the long axis of the body, but this is not the same thing as up and down, which continue to be defined in terms of the gravitational field, even if one is leaning at an angle. The distinction between left and right does not follow either of these models. The fundamental distinction here does not have anything especially to do with the bodily axes at all. It is not, as in the case of in front and behind, that there is any bodily asymmetry that the distinction labels, since animals are generally right-left symmetric. Nor does it label, like 'up' and 'down', some external physical magnitude of general importance for action. Nonetheless, it is evidently an axis used to direct action.

I said that egocentric axes are "immediately" used to direct action. It may be that no very precise definition can be given of this notion of immediate use, and that the notion of an egocentric reference frame must to this extent remain a rough

and intuitive one. But we can get some sense of the required conception by contrasting egocentric frames with more complex dead-reckoning systems, that is, systems that enable one to keep track of where one is by keeping track of how fast one has been moving, in what direction, and for how long. The point about such systems that matters here is their use of a compass, which is external to anything used to immediately direct action. For example, an animal might use the position of the sun, together with its knowledge of the time of day, as a compass. It can use the position of the sun to keep track of each of its various swoops and sallies, and so to plot the direct route home. But before it can actually translate this into action, it has to know the direct route home not merely in terms of direction specified in terms of the external compass; it has to know which way to point itself to travel in that direction. It is in this sense that the egocentric axes are immediately used to direct action, whereas the external compass is not. Of course, a dead-reckoning system could also use the egocentric axes themselves, though in practice this would mean a considerable loss in accuracy. Notice, incidentally, that these dead-reckoning systems are body centered; what make them nonegocentric are the axes that they use and the indirectness of their role in guiding action.

So far we have considered one part of the characterization of a frame of reference: the way in which it identifies places at a time. But characterization of a frame of reference must also say what criterion is being used for the identity of places over time. To pinpoint the issues, consider an extremely simple experimental paradigm used by Linda Acredolo (1990) to find out whether and in what way infants identify places. Her experi-

mental space is an enclosure ten feet square with two identical windows across from each other. There is a round table in the center of the room with a buzzer under it and a long moveable rod attached. And the end of the rod is a seat on wheels, which can revolve around the table. On top of the seat is an infant. In the training phase, the buzzer sounds in the center. About five seconds later, an adult appears at one of the windows, calling the child's name and generally entertaining it for five seconds or so. Of course, the child turns to look, and the pairings of buzzer and event, always at the same window, continue until the child has developed an expectation that such an event will follow the buzzer. That is, on hearing the buzzer, the child turns toward the window before the adult appears. After this training phase, the chair is moved around the table to the other side of the room. The buzzer is sounded, and the experimenters watch to see toward which window the child looks in expectation of the event.

Obviously, if the child has learnt merely a spatial response, such as to look to the left, it will look toward the wrong window. There certainly are these spatial behaviors, though they are more primitive than the ability to identify places. For example, there is the ability to reach to the left or the right, or to jump out of the way of an oncoming object. Even if the infant has only a particular response, such as looking to the left, its behavior may still be properly described as spatial. It may vary the type of muscular movement in many different ways, depending on the starting orientation of its body when the buzzer is sounded, so as always to achieve the result of looking to the left. So it may be impossible to describe the response as a nonspatial muscular movement, even though it is more primitive than place identification.

It might indeed be said that in the case in which the child looks at the wrong window, it is using a notion of place on which sameness of place over time is defined by sameness of egocentric coordinates. On this notion of place, no sense can be assigned to the idea that the child itself might be in motion, or capable of movement. Rather, it has an array of places, such as the one "just within reach and to the right," that it carries with it through the world. In this frame, something is said to be in the same place at one moment as at another if at both times it was just within the subject's reach and to the right, whether or not the subject had, as an observer using a more standard frame of reference might say, "moved" in the meantime. We can certainly imagine a subject for whom this is a possibility. For example, an Oriental despot might so arrange matters that however and whenever he moves, there is always Turkish delight just within reach and to the right.

We have finally to consider the case in which the child manages to use the information available to it through the rotation to keep track of the right window and look toward it, even though this means giving a different spatial response, such as looking to the right rather than looking to the left. In this case the child is certainly reidentifying a place, but it is not using sameness of egocentric coordinates over time as the criterion for sameness of place. This raises the question what criterion the child is using. The obvious proposal is that the criterion used is egocentric coordinates plus compensation for the child's own movements. No frame of reference other than the egocentric frame is used.

Acredolo found a gradual transition from looking at the wrong window to looking at the correct window as the children grew older or, more precisely, as the time during which the child was capable of self-locomotion increased. Similar

keeping track of a place over time is achieved by the whole class of animals that manage to find their way back home by keeping track of their own movements—the directions and distances of their travel from moment to moment—and then using path integration to find the direct route home (see, for example, Müller and Wehner 1988).

1.3 Maps

As we saw, some people have denied that there is such a thing as objective or absolute space. Empiricism or pragmatism about spatial concepts means that a spatial frame of reference must always be understood in terms of its relation to egocentric space. Poincaré put the point succinctly: “absolute space is nonsense, and it is necessary for us to begin by referring space to a system of axes invariably bound to our body” (1913, 257). This was vigorously rejected by O’Keefe and Nadel:

Most authors attempt to derive all psychological notions of space from an organism’s interaction with objects and their relations. The notion of an absolute spatial framework, if it exists at all, is held by these authors to derive from prior concepts of relative space, built up in the course of an organism’s interaction with objects or with sensations correlated with objects.

In contrast to this view, we think that the concept of absolute space is primary and that its elaboration does not depend upon prior notions of relative space. . . . [There] are spaces centred on the eye, the head, and the body, all of which can be subsumed under the heading of *egocentric space*. In addition, there exists at least one neural system which provides the basis for an integrated model of the environment. This system underlies the notion of absolute, unitary space, which is a non-centred stationary framework through which the organism and its egocentric spaces move. (1978, 1–2)

Let us consider what the prospects are for finding such an absolute or objective mode of spatial thought. We saw that we

cannot define an egocentric space as a body-centered one, but can we not define an allocentric frame of reference to be one centered on something other than the body? Certainly that is in practice often what is meant by 'allocentric', which is a term often opposed to 'egocentric' in the literature.

There are many distinctions to be drawn here. For example, we can ask whether it is possible for a subject to be using only allocentric frames of reference in spatial action. In the case of a patient described by Cole (1991) who has no or only residual proprioception and kinesthesia, for example, the patient describes the extraordinary effort involved in purely visually guided action: taking a visual fix on a point in the room and keeping himself upright and stationary only by maintaining his fix on the point; picking up a glass by remarking its relation to his fixed point, the relation of his hand to the fixed point, and bringing the hand to the glass by varying its relation to the fixed point. It does not seem to be an unintelligible hypothesis that in action the subject may be using only an allocentric frame of reference, one centered on his fixed point, and that there is no immediate use of any body-centered frame of reference. In that case, an allocentric frame of reference would be what is immediately used in guiding action. The allocentric frame would then count as "egocentric" by the definition I reached in section 1.2. This is discordant terminology. The fact is that I am drawing distinctions that the literature obliterates; what matters is that we separate and understand the various classifications that there are. The use of 'egocentric' as defined above is also widespread. The point is that 'body-centred' and 'immediately used in directing action' ordinarily coincide, which is why ordinary vocabulary does not distinguish them.

Would the existence of an allocentric frame of reference, one not centered on the subject and that need not be immediately

used to guide action, show that there is an objective level of thought that resists empiricist-pragmatist criticism of the Poincaré type? I can make the question concrete by considering a very simple representative of a whole class of navigational systems (see Cartwright and Collett 1983). This is the triangulation model used by Wilkie and Palfrey (1987) to explain the behavior of rats in a water maze. 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 track of 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 platform goal, the animal swims toward 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. The animal using this model certainly has an allocentric representation of spatial relations, in the sense that it uses a frame of reference centered on something other than its own body. It has recorded the distances from the target platform to the cues around it. But it is hard to see why the empiricist or pragmatist should be particularly moved by this. For it is only through its connections to its own perceptions and actions that the animal manages to give any meaning to the spatial information. The allocentric representation has meaning only through its connections with the animal's egocentric space, in the sense of 'egocentric' defined above.

There is another array of distinctions that I should remark on. We already saw that for a creature to be representing places, it must have some grasp of a criterion of identity for places over time. It is a separate question whether it appreciates that the places it represents are all spatially related to one another, whether it grasps the connectedness of the space.

Consider again the child who succeeds in Acredolo's (1990) paradigm. The child is certainly representing places and may be able to represent more than one place at a time, but this does not yet show that it is capable of grasping the spatial relations among the places it can represent.

It would be possible to have a creature or system that simultaneously made explicit all the spatial relations between all the places it could represent. This is a very strong condition. It is also possible to have a creature that, while not meeting this condition, can make explicit a sufficiently rich range of spatial relations and perform sufficiently powerful operations on that base of spatial relations to derive the spatial relations between any two places.

There is a family of conditions here, depending on exactly which spatial relations we have in mind. There can be variations in what the underlying geometry is taken to be. And there are differences in what configurational properties of a network of places the system might be capable of representing.

These issues about the extent to which the connectedness of a space is represented are just different from questions about where the frame of reference is centered. It would be possible to have a system capable of very powerful representation of connectedness centered on the body, for example.

Possession of a map of one's environment is sometimes defined in terms of the strength of the spatial relations that can be represented among the places represented.¹ Mapping abili-

ties in this sense constitute no very evident objection to the empiricist-pragmatist critique of the notion of objective space. I can bring this out by turning to the model of spatial navigation proposed by John O'Keefe.

I will describe O'Keefe's (1990, 1991) model only very schematically. On this model, the *slope-centroid model*, there are two stages in an animal's construction of a map of its environment. The animal identifies a notional point in its environment, the centroid, which is a notional point in the sense in which the South Pole or the Equator are notional: there may be no distinctive physical feature at that place. It is a fixed point, in that it does not move with the animal. The animal also identifies a gradient for its environment, a way of giving compass directions. This is the slope of the environment. It functions like the direction east-west. The direction is fixed no matter how one moves around, and one can partially define which way one is going by saying what angle one is making with it. As in almost all models of mapping, we take it that the animal is constructing a two-dimensional map of its environment; the third dimension is not mapped.

Once the animal has done this, it can construct a map of its environment by recording the vector from the centroid to each of its targets, using the slope to define direction. Suppose that it has done this and now wants to know how to get to a particular target. What it must do is to find the vector from itself to the centroid. Once it has the vector from itself to the centroid and the vector from the centroid to the target, it can find the vector from itself directly to the target.

This certainly gives the animal an allocentric frame of reference, in the sense defined above, for the frame of reference is organized around the centroid, rather than the body of the

animal. This is also a geometrically very powerful model, capable of representing many of the spatial relations among the places in its environment. So it will meet any reasonable connectedness condition on maps.

The question that remains is whether the existence and use of such a system constitutes a reply to the empiricist-pragmatist critique of the notion of objective space. Does it fare any better than the triangulation model? We saw that from this point of view, the problem with the triangulation model is that it seems to have meaning for the animal only insofar as it is connected to the animal's perception and action, to its egocentric space. Now the slope-centroid model does seem to be more distanced from perception and action than the triangulation model. The reason is the purely notional character of the slope and centroid, which do not themselves relate directly to perception: they are computed on the basis of it. In contrast, the distinguished point in the triangulation model is the platform goal, which the animal actually occupies and from which it observes the distances to the cues around it. Yet despite the fact that there is this sense in which the slope-centroid model seems to be more remote from perception and action than the triangulation model, it still seems that the model has meaning for the animal only insofar as it is connected to perception and action. The basic point here is quite elementary. The mapping systems we are considering are all navigational systems. Their significance is exhausted by their implications for navigation. It is, therefore, quite impossible that they should constitute objections to the empiricist-pragmatist critique of objective space. If we want to find examples of genuinely objective spatial thought, we have to look for modes of spatial thought whose role lies not only in the demands of navigation.

1.4 Physical Objects and Objective Space

The fundamental point in all this is that we cannot ascribe spatial representations to animals in a way that outruns their capacity to give causal significance to the representations. And the distinctions that matter are distinctions in the way causal significance is assigned. I want to elaborate on this point before returning to the characterization of objective space.

We can distinguish between a pure geometry, which is a purely formal exercise in mathematical computation, and an applied geometry, which is a body of doctrine about the world in which we live. What turns one into the other is the assignation of some physical meaning to the spatial concepts, for example, the identification of a straight line as the path of a light ray *in vacuo*. So, to find what about spatial reasoning makes it reasoning about the space one is in, we must look for an account of the laws or regularities, however probabilistic or open to exception, that connect spatial properties with other physical properties. To do this is not to demand definitions of spatial concepts in terms of other physical concepts. We cannot assume that spatial concepts are less primitive than any others, and we cannot assume that we can separate out, from the flux of physical thinking, some regularities that alone deserve to be elevated into definitions. There may be no definitions to be had (M. Friedman 1983). Still, insofar as spatial reasoning is to be understood as reasoning about one's actual environment, rather than as pure geometry, it is theoretical. It is only its figuring in an intuitive physics of one's environment, through regularities connecting spatial properties with other physical properties, that makes it reasoning that is not purely mathematical but rather about the space in which one lives.

We also have to explain why an animal's capacity to engage in spatial reasoning might have been selected for under evolutionary pressure. And if the reasoning has no physical significance, it cannot help the animal to cope with its surroundings and thus has no selection value. We ought to be reluctant to ascribe spatial reasoning when its use by the organism defies explanation in terms of selection pressures. If we subscribe to an evolutionary-teleological view of content—*ascription* (as in Millikan 1984), we will not simply be reluctant to do this but regard it as incoherent to do so. Of course, there are views on which spatial reasoning is simply a form of causal reasoning: spatial notions can in some sense be reduced to, or explained in terms of, causal notions (Sklar 1983, van Fraassen 1985). But we can accept the need for us to relate a creature's ability to represent space, to its capacity to give physical significance to those spatial representations, without insisting on any such reductionist thesis. We can even accept that spatial notions have to be used in elucidating the concept of cause (Salmon 1984). It seems unlikely that there will be any firm formal rule governing the connection between the spatial content of reasoning and the physical interpretation assigned to such content, just as in physics there is no firm formal rule governing the physical meaning that a theory must assign to a spatial magnitude. All we can do is to look at the ways in which particular theories use particular spatial magnitudes and, case by case, argue that this or that one has or has not been assigned sufficient physical meaning.

Suppose, now, we ask what use one makes of grasp of the fact that every place in one's environment is spatially related to every other place in it. The answer is that we have to look at the physical significance one assigns to this relatedness, at the way one grasps the causal connectedness of space. Because of

what I will call their *internal causal connectedness*, physical objects play a special role in how we register the connectedness of space. The internal causal connectedness of objects means that the possibility of their traveling through a space can give physical significance to its spatial connectedness.

What does it mean to say that physical objects are internally causally connected? The point is that the condition of a thing at any one time is causally dependent upon its condition at earlier times. Grasp of this idea is presupposed in an understanding of the way in which objects interact with one another. If we are to have any appreciation at all of the effect that one object can have on another in a collision, for example, we have to understand that one determinant of the way the thing will be after the collision is the way that very thing was before the collision. The result of the impact may be a smash or a bounce; which result happens will depend in part on how brittle the object was to begin with. The way the object is later depends in part on the way it was earlier, and we have to grasp this if we are further to grasp that the earlier condition of an object is only a partial determinant of the way it is now and that external factors may have played a role. So in describing our ordinary thought about physical objects, we need a distinction between the causality that is internal to the object and has to do with the dependence of its later stages on its earlier ones and the causality that has to do with the external relations between objects and the ways in which they act upon each other.

Even if we consider a physical object through a period in which it is not involved in any interactions, it remains true that its condition through any interval in that period causally depends on its condition in earlier intervals. This is so whether the condition of the object remains stable during the period or is inherently subject to some variation, such as a decay process.

In this matter, objects can be contrasted with, for example, shadows. Even if a shadow remains constant through a period or undergoes only regular variation, its condition at times through that interval does not causally depend on its condition at earlier times. Rather, the condition of the shadow at any time depends directly on the way things are with the light source, occluders and surfaces. It does not depend on how things were with the shadow earlier. So unlike physical objects, shadows are not internally causally connected. The same point could be made about, for example, the spot of light cast by a searchlight on a wall, which is like a shadow in not being internally causally connected.

Philosophers interested in causation have tried to use this point about physical objects in analyzing causation; philosophers interested in the identities of physical things have tried to use this point in analyzing object identity.² But without attempting to analyze either causation or object identity, I can say what is the bearing of this point on the way in which we give causal significance to the spatial connectedness of a region. The internal causal connectedness of physical things means that they can give physical meaning to spatial connectedness. In particular, the possibility of an object's moving from one place to another means that we can see how the way things are at one place could causally depend on the way they were at another place. For example, if a horse plods from the start of a track to its end, the way things are at the end of the track now causally depends upon the way they were at the start. When we consider the movements of objects through a space as causally connecting one place with another, if we further consider the details of that movement, we can see how to give physical significance to the metric for the space within our intuitive physics. The crucial notion here is the time taken for the object to

reach a particular destination from a particular starting point, given what sort of thing it is and what causes are affecting its movement. These remarks only begin to sketch the structure given to the space of our intuitive physics by physical things. The continuity of object movement means that an order is imposed on the places between starting point and destination, depending on the trajectory we ascribe to the object, and this order in turn is responsible to our conception of the causes of the movement of the object. None of this can be achieved by considering shadows alone. Just because they are not internally causally connected, shadows cannot be used to give causal significance to the spatial connectedness of a region. The movement of a shadow from one place to another is not a way of ensuring that the way things are at the destination is causally dependent on the way things were at the starting point.

So much for the role of thought about physical things in giving causal significance to the spatial connectedness of a region. I want now to remark that there seems to be a level of thought more primitive than our thought about physical objects, and to consider how connectedness might be thought about at that level. 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, we may have 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 this. So it may be an effort to recall astronomical knowledge from school. That is, I think, part of the reason why men landing on the moon was such a shock. It is not just the distances involved. It

forced us very directly to think of the moon as a physical thing, and this is not ordinarily forced so vividly upon us.

Consider now the case of an animal swimming in the water maze, using the triangulation model and the cues hung around the pool to navigate to the platform goal. 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 this. They are recognizable and stably at those places, and this is really all the animal needs.

Nevertheless, as we saw, the animal using the triangulation model does manage to give causal significance to the spatial relations between various places in the space it is in, even if it does not manage to register the full connectedness of the space. But evidently, it cannot be doing this in the way I have just been describing, exploiting the internal causal connectedness of physical objects, since it may not be capable of thinking in terms of physical objects at all but rather be operating at a more primitive level of thought. So how is it managing to give causal significance to spatial notions? It does not do this by reflective causal thinking at all. It does this through its own engagement in the space: not by thinking about its engagement in the space but by putting the triangulation model to work in navigating itself through the space. I could talk about a *practical* interpretation of the spatial notions used in the triangulation. This makes very obvious why thought at the level of the triangulation model cannot be described as objective or absolute spatial thinking. For such thought gets its physical significance only through its relation to the subject's perception and action, only through its relation to his egocentric space.

Let us now consider the model of slope and centroid. What motivates the idea that this is a peculiarly objective or absolute type of thought? What makes this seem plausible is that the vectors from the centroid to the cues in the environment do not depend on the current location of the animal. The animal can record these vectors and carry that map around with it. This is why we seem to have a level of thinking here that does not have to be explained in terms of its relations to the animal's perception and action, to its egocentric space.

It is striking, however, that if we ask whether the animal using the model of slope and centroid must be thinking of the cues as physical objects, the answer is no. It may be using the primitive level of thought that we can use when thinking about the stars, at which the distinction between objects and shadows is not yet drawn. It need not be thinking of the cues as internally causally connected. It may, for example, have no expectations whatever as to what would happen if two of the cues collided. The way in which causal significance is being assigned to the spatial notions is not through the thought of the things around it as physical objects.

There is a contrast between the cues that an 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 may 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 have some persisting effect on it. It need not have a reflective understanding of these

points, but it does have to have some practical grasp of them. The reason that it does not need a reflective grasp of these points is that it is using its own interactions with the targets to constitute its grasp of their causal significance.

Certainly the animal is not using reflective thought about its targets to assign causal significance to spatial relations: it does this rather through the fact of its own engagement in the space. The slope-centroid model is in this regard in exactly the same position as the triangulation model. It is true that the vectors from the centroid to various cues do not depend upon the location of the animal. But when we subtract the animal, we also subtract any physical meaning for those vectors. The vectors get their meaning only when the animal is plugged in. They have causal significance only through their relations to the animal's perception and action, to its egocentric space.

At this point it is not hard to see what is required of an objective or absolute representation of space. What is right about the empiricist-pragmatist critique is that spatial notions must be given causal significance. The mistake is to think that this can be done only through one's own interactions in space. What we need is a way of registering the connectedness of a space in a way that does not depend upon the subject's own engagement in the space. And we have seen why physical objects might be expected to play a crucial role here. There are, of course, many other phenomena than the movements of physical things that in diverse ways transmit the effects of things being thus and so at one place to their being thus and so at another place. For example, there are the everyday phenomena of magnetism, heat and cold, the flow of liquids, and the winds. One fundamental range of alternatives to physical objects emerges if we consider a mariner navigating in a vast circuit of tides, whirlpools, eddies, and currents. It is open to him,

in principle at any rate, to register the physical significance of the spatial connectedness of the region he is in without exploiting his own navigation through the space or introducing the notion of a physical object. The waves themselves, propagated through the space and interacting with one another in endlessly complex ways, demand for their understanding a rich grasp of the connectedness of the space. In our commonsense understanding of the world, we can, to some extent, use this kind of causal thinking on land, as when we watch the effects of an earthquake or the impact of a sledge hammer on the wall of a house. But these phenomena are not sufficiently pervasive in our experience to provide the full strength of our grasp of the theoretical significance of the connectedness of the space we occupy. What we have to investigate now, and what I will begin on in the next chapter, is the way in which one's conception of one's life as extending over time, a time in which one has a current temporal orientation, depends on this ability to form an objective conception of the space one is in. This is what explains the value of objective spatial thinking for us.

1.5 Common Causes and Informative Identities

In *Individuals* (1959), Strawson held that reidentification of places depends on the reidentification of things. We identify places by their relations to things; we see that we are once again between Black Mountain and the lake, for example. In view of the dependence of the reidentification of places on the reidentification of things, Strawson said, "The fact that material bodies are the basic particulars in our scheme can be deduced from the fact that our scheme is of a certain kind, viz. the scheme of a unified spatiotemporal system of one temporal and three spatial dimensions."³ This thesis is certainly too

strong as it stands. An animal might use a set of cues to identify and reidentify places without ever raising the question of whether they are processes or pseudoprocesses; the animal might operate at a level of thought that is more primitive than that at which a distinction is marked between a shadow and a thing, since a stable shadow will do as well as a thing when reidentifying places.

Nevertheless, it can hardly be denied that our reference to places is densely interwoven with reference to things and that reference to things greatly enhances our capacity for reference to places. This is principally because of the possibility of perception-based informative identities of physical things. On recognizing the identity of Mount Afla with Mount Ateb, for example, one opens up the possibility of relating all the places one had identified using Afla to all the places one had identified using Ateb. One can now grasp the spatial relations of all these places to one another. Again, it can happen that, lost in a town one has just come to, one suddenly realizes that one is staring at an unfamiliar angle of a building one knows perfectly well and so can find one's way from there.

To understand such happenings, we have to bear in mind a further causal dimension to physical objects other than that they are internally causally connected. This is the capacity of a physical thing to function as a common cause of many different effects. To make the point vivid, I will move away from cases of ordinary navigation and look instead at an astronomical example. Consider the case of the twin quasars 0975 + 561 A and B. These objects are close together in space, much closer than seems likely, given the scatter of the other known quasars, and their spectra are indistinguishable. The similarity here is rather unlikely, and it has an obvious explanation. Namely, we

should assume that the minor discrepancies in the information we have about the two quasars are due simply to the margin of error in our methods and that the quasars actually have all their properties in common; that there is just one quasar there. This kind of reasoning is absolutely standard in cases in which an informative identity is discovered. We find that two things are highly similar; we explain the similarity by postulating an identity of objects and remarking that by Leibniz's law, they therefore have all their properties in common.

I must also remark, though, that there is another way of reading the case of the twin quasars. We have two images of the quasars with an astonishing correlation in the information they contain about the quasar that is their source. This correlation we explain by the fact that there is a common cause of both images. Since the same thing causes the production of both, the correlation is intelligible. In the case of the quasars, what happens is that a galaxy intervening between the quasar and us splits the light, and consequently the image. These are two quite different descriptions of the same example. The problem is not to choose between them, for surely both are illuminating. The problem is to understand how they can both be correct and to understand the relation between them.

What are the effects that the quasar is being said to be a common cause of? The obvious answer is the production of two images of itself at some spatial separation. The similarity between the two images, then, is what requires explanation in terms of a single common cause. Observe, though, that it is not crucial to the example that there should be an intrinsic similarity between the images themselves; the similarity that matters in the reasoning that leads one to the conclusion of identity is a similarity between the quasars represented. To suppose that

this similarity can be represented only if we have some intrinsic similarity between the images themselves is just a mistake. Representation of similarity need not imply similarity of representation (Millikan 1991). The principle of the common cause is being invoked to explain why there is the otherwise improbable correlation between the content of the image of the one quasar and the content of the image of its twin.

The case of the quasar is somewhat unusual among cases of informative identities in that there is explicit reflection on the character of the images formed and the resemblances between them. In the case of the Morning Star and the Evening Star, there may be no such reflection on our perceptual images. We may simply remark, on the basis of our perception of the sky, that the Morning Star and the Evening Star have much in common and then appeal to identity and Leibniz's law to explain the commonality. But even so, the principle of the common cause is still in play. For the properties that we note the two bodies to have in common must be causally significant. And the identity statement, that the Morning Star is the Evening Star, is not true unless the cause of one's perception of the one heavenly body is the same as the cause of one's perception of the other heavenly body. There is, then, a sense in which the reasoning behind an ordinary informative identity in which there is no reflection on the perceptions of the object nevertheless depends upon the principle of the common cause. The ground-floor reasoning is correct only if the reflective recapitulation, which refers to the perceptions and explicitly uses the principle of the common cause, also correct.