

The Nature of the Senses

Matthew Nudds

University College London

PhD

September 1999

Abstract

My thesis provides an account of the nature of the senses. Many philosophers have supposed that the fact that we have different senses makes the integration of the senses problematic. In this thesis I argue that introspection reveals our perceptual experience to be amodal or unitary (that is, we cannot distinguish distinct experiences associated with each of our senses) and hence that the real problem is not how the senses are integrated with one another, but how and why we distinguish five senses in the first place. What we need is an account of what our judgements are about when we judge that we are, say, seeing something or some property.

I argue that such an account cannot take any of the forms commonly supposed. Philosophers often assume that an account must appeal to differences between kinds of experience, but I argue that such differences are not sufficient to explain the way that we distinguish five senses. Nor can we explain the distinction by appealing to the different kinds of mechanism involved in perceiving, since recent cognitive psychological models of the mechanisms of perception show them to be functionally diverse in a way that undermines any correspondence between them and the five senses, and our common-sense grasp of the different mechanisms involved in perception presupposes a prior understanding of the distinction between different senses.

I provide an account of the distinction that we make between the five senses, according to which the senses are not substantially distinct. Although our judgements about the senses are true, they are not judgements about kinds of thing; rather, we distinguish different ways of perceiving in terms of different, conventionally determined, kinds of perceptual interaction we can have with our environment.

Contents

1. INTRODUCTION.....	4
2. THE UNITY OF EXPERIENCE.....	17
3. THE PRODUCTION OF SOUNDS	37
4. THE FEATURES ACCOUNT.....	59
5. THE BRUTE EXPERIENTIAL VIEW	84
6. A SIMPLE THEORY OF PERCEPTION	105
7. THE SENSES AS NATURAL KINDS.....	128
8. THE SENSES AS CONVENTIONAL.....	163
BIBLIOGRAPHY	192

1. Introduction

Intuitively, seeing and touching something triangular are of distinct and conceivably separate experience types (Loar 1996, p.321).

We all of us, or at least nearly all of us, distinguish five different senses. That is, we all have names for the five senses of sight, touch, hearing, smell and taste; and five corresponding verbs of perception. We make judgements involving concepts of just these senses, both about the way that we ourselves perceive things, and about the way other people perceive things. We judge, for example, that we can see the blackboard, that we can hear the music, that the table feels smooth, and we judge that someone is smelling a flower, can't see the river, or that they heard what was said.

Distinguishing five senses in this way appears to be almost universal,¹ and it is not specific to the English language. Nor is it a recently invented distinction or an artefact of our scientific culture (in the way that our concepts of, for example, a film or photograph perhaps are). Aristotle theorised about the five senses more than two thousand years ago, and devoted a large part of his treatise *De Anima* to a discussion of them.² But Aristotle didn't invent the five senses and he wasn't the first person to make the distinction, nor the first to theorise about the senses. A Greek philosopher-cum-physician, living perhaps two or three hundred years before Aristotle, is reported by Theophrastus to have "discussed each of the senses," although he doesn't seem to have got very far,

he says we hear with our ears...smell with our noses...discriminate flavours with our tongues...[that] the eyes see...[and that] all the senses are somehow connected to the brain. As for touch, he said

¹ The practice of distinguishing different senses appears to be universal, that of distinguishing five may not be. See Ritchie (1991), and Chapter 8 below.

² See Hamlyn (1975). For a discussion of Aristotle's theory of perception and the senses see Everson (1997) and Modrak (1987).

neither how nor by what means it works. So much for Alcmaeon's views (Theophrastus 1987).

It seems that these writers were not inventing or uncovering a distinction, but theorising about something that was already for them part of everyday experience. Indeed, that we have five senses can just seem an obvious and undeniable fact about us. But what are the senses? What are our concepts concepts of? This is a question that has rarely been directly addressed by philosophers. Those who have addressed the question disagree about how to answer it,³ and none of their answers (which I discuss throughout this thesis) are satisfactory. Of course, much recent philosophy has been concerned in one way or another with our perceptual experience of the world, but this recent discussion has tended to concentrate, implicitly at least, on vision and visual experience. (So much so that it can sometimes seem as though philosophers have forgotten that not all perception is visual perception.) It is assumed that whatever account we give of visual experience will simply generalise to the other senses. It is assumed, too, that an account of the senses will simply fall out of an account of perceptual experience in general. In this thesis I shall argue that once we get clear about what an account of the senses must explain, the latter assumption will be seen to be unfounded: although it is obvious that we have five senses, what the five senses are is far from obvious.⁴

Whenever we find it difficult to account for a common-sense concept we may draw the sceptical conclusion that there is no account to be given. We may conclude, for example, that the concept is illegitimate, that we cannot give any coherent account of the conditions which govern its application. Alternatively, we

³ Grice's paper (1967) is still the best thing to have been written about the senses; Coady's (1974) is a comment on it. See also Roxbee-Cox (1970), Sorabji (1971) and Sanford (1976), and more recently Leon (1988) and Nelkin (1990).

⁴ When talking about the five senses I mean the five senses of common sense; it is our common sense or everyday use of these concepts that I want to explain. There is a notion of 'sense' which perhaps has some scientific connotation and is different from that of the five senses. It is this different notion of sense that is in use when people suggest that we have more than five senses; given that notion of a sense such a claim might be true. But that is not the notion of a sense that I am giving an account of.

might adopt an error theory of the concept in question.⁵ To adopt an error theory of some concept is to accept that the concept has determinate conditions of application, and so to accept that judgements involving that concept have truth-conditions, but to deny that there is in fact anything to which the concept applies, and so to deny that judgements involving it are ever true. *Prima facie*, there's no reason to think that our concepts of the senses are illegitimate, but it is not necessarily absurd to suggest that we might adopt an error theory of them. It might be, for example, that although we do have different senses (revealed, perhaps, by scientific investigation), we don't have the five senses that we commonly think we have, and hence that our modality judgements are all false. It might then be argued that we should eliminate our common-sense concepts of the five senses and replace them with more scientifically respectable concepts.⁶ Whatever we should say in general about these kinds of sceptical position – and I think that most people would accept that the global elimination of folk-psychological concepts is a non-starter – it will be an assumption throughout this thesis that we should not adopt an error theory of our concepts of the five senses. I think this assumption justified for two reasons. The first is simply a prejudice. When a concept is part of our common-sense or folk-psychological conceptual repertoire – in the way that our concepts of the senses are – it must play a role in our lives; were it not to do so (in fact, were it not effectual in doing so) it would not have remained part of common sense. The very fact that it is part of common sense psychology gives us a reason for thinking it has application. So if we can't explain such concepts as being concepts of some, say, scientific kind, we shouldn't conclude that common-sense understanding is somehow mistaken, but rather that these concepts are to be understood in some different way.⁷ (In general, I think it mistaken to assume that scientific psychology is simply a more extensive and precise development of folk psychology.)

⁵ See Mackie (1977, p.35) for the idea of an error theory.

⁶ For eliminativism about propositional attitudes (and folk-psychological concepts in general) see Churchland (1981). For an argument that our folk-psychological concepts of emotions may be mistaken see Griffiths (1997).

⁷ This is not true in general of our folk concepts, and may well not be true of our folk-physical concepts, for example.

The second, and decisive reason for rejecting an error theory, is that we can give an account of our concepts of the five senses which explains both how they have application and how judgements involving them can be true. It may turn out, according to such an account, that the five senses are not quite what we thought them to be, but not that our concepts lack application. In this thesis I give such an account. Although I don't answer all the questions about the senses that need to be answered, nor solve all the problems to do with the senses that need to be solved, I do give an account of what the senses are, and so provide a framework within which these further questions can be answered and problems solved.

1. That we have five senses seems obvious, yet those who have attempted to give some account of them disagree fundamentally about their nature. In giving an account of the senses, what is it that we must give an account of? What is it that we need to explain? Verbs like 'look', 'feel', 'taste', and 'smell' "are very frequently used to refer to publicly observable operations by persons," including ourselves (Coady 1974, p.112). We talk, for example, of someone smelling a rose, or looking at a picture, or listening to a piece of music, and so on. We need to explain on what basis and why we apply concepts of the senses to various activities that we and others engage in, and what makes these activities the activities they are. We also use the concepts of the senses to refer to the psychological effects of these activities – to various kinds of experience. We talk, for example, of things looking red, or of the rose smelling fragrant, and so on. We need to explain, then, why we apply concepts of the senses to such experiences.

The way that something looks is different to the way that it feels, which is in turn different to the way that it sounds. (To confirm this think, for example, of the way a rose looks compared to the way that it feels.) Since we can see something without feeling it, and feel it without seeing it, the way that something looks is, *prima facie*, distinct from, and independent of, the way that it feels. Intuitively, then, there are different and distinct types of experience associated with these two senses, indeed with each of the senses – there are visual experiences, tactual experiences, auditory experiences, and so on. If that is correct, then an account of the senses needs to explain what makes these experiences the kinds of experience they are; that is, what is distinctive of the kind of experience associated with each of the senses.

If we can provide such an account, then we can use it to explain why we apply concepts of the senses to the various kinds of activities people engage in: they are the activities which tend to produce these kinds of experience.

I suggested that an account of the senses needs to be an account of the distinct kinds of experience that, intuitively, we take to be associated with each of the different senses. Many writers on perception share these intuitions: they suppose that there are distinct kinds of experience associated with each of the senses and, whether or not they think that an account of the senses must explain anything more, they suppose that an account of perceptual experience must at least be able to explain what is distinctive of these different and distinct kinds of experience. This requirement is, however, not usually taken to place any distinctive constraint on one's account of the nature of perceptual experience. Rather, one's account of the nature of experience is determined by one's attitude to various other more general metaphysical and epistemological problems – in particular one's attitude to the argument from illusion. An account of the distinctive kinds of experience associated with each of the senses can be expected to be determined by the more general account of the nature of experience formulated in response to these problems.

So, for example, a sense datum theorist takes our perceptual experience to consist in our awareness of, or acquaintance with, certain kinds of particular objects – sense data – which are taken to be mind-dependent by some theorists, as non-material by others.⁸ Two experiences are distinct, according to this kind of view, if they are constituted by the awareness of distinct sense data. We can distinguish kinds of experiences in terms of the awareness of sense-data. If we can group the objects of experience – the sense data – into kinds, so too can we group the experiences themselves – the awareness of such sense data – into corresponding kinds. Two experiences are of different kinds if they are instances of the awareness of different kinds of sense data. The sense-datum theorist can, therefore, explain what makes an experience the exercise of a particular sense by appealing to the

⁸ Some sense-datum theorists are non-committal about the status of sense-data; see, for example, Moore in (1918-19). He later changed his mind (1957). For sense-datum accounts see also Price (1932), Jackson (1977), Perkins (1983), and Robinson (1994).

kinds of sense-data awareness of which is constitutive of them. H.H. Price, for example, is typical when he says that

there are not different kinds of sensing. Visual sensing will simply be the acquaintance with colour-patches, auditory sensing the acquaintance with sounds, and so on; the acquaintance being the same in each case (1932, p.5).

Berkeley's view of experience is slightly different. He thinks that our perceptual experience is constituted by an awareness of instances of qualities – what he calls sensible ideas – and that there are different kinds of quality specific to each of the senses. We don't, he says,

immediately perceive by sight any thing beside light, and colours, and figures: or by hearing, any thing but sounds: by the palate, any thing beside tastes: by the smell, beside odours: or by the touch, more than tangible qualities (1734, p.175).

Particular experiences of different senses are distinct because they are constituted by an awareness of distinct instances of sensible qualities. The experiences associated with each sense are of different kinds because each is constituted by the awareness of a different kind of sensible quality:

the extension, figures and motions perceived by sight are specifically distinct from the ideas of touch, called by the same names, nor is there one idea, or kind of idea, common to both senses (Berkeley 1732, section 127).

Many recent theorists of perception have rejected a sense-datum or Berkeleian theory of perception and have embraced instead an intentional or representational

theory of perception.⁹ According to this kind of view, we should think of our experiences as intentional states which represent things in the world around us as being a certain way. What it's like for one to have an experience is a matter of how that experience represents things as being. Although such theorists view the nature of experience differently to sense-datum theorists, they still suppose that an account of the senses must explain what is constitutive of the different kinds of experience associated with each of the senses. Dretske, for example, holds what he calls the Representational Thesis, the thesis that all mental facts are representational facts, from which it follows that

the quality of experience, how things seem to us at the sensory level, is constituted by the properties things are represented as having. My experience of an object is the totality of ways that object appears to me, and the way an object appears to me is the way my senses represent it (1995, p.1).

He then goes on to explain what the difference between the experiences of different senses consists in. According to Dretske,

[a] blind person may know what it is like to visually experience movement. If he knows what movement is, that is enough. An experience of movement – whether it be visual, tactual, or kinaesthetic – has its qualitative character defined by what it is an experience (representation) of, and if these experiences are all of the same property, they are, subjectively, with respect to this single property, the same kind of experience. In knowing what it is like to visually experience movement, though, a blind person does not necessarily know what it is like to visually experience an object that is moving. For there is more – much more – involved in seeing an object move than experiencing the object's movement. One also

⁹ See, for example, Harman (1990), Peacocke (1992, Ch.3), and Dretske (1995). There are others; the view is, one might say, the current orthodoxy.

experiences the object's shape, size, color, direction of movement, and a host of other properties. This is why seeing and feeling movement are much different even though the same thing (movement) is represented in both modalities. Even when the senses overlap in their representational efforts – as they do in the case of spatial properties – they represent different ranges of determinable properties (1995, pp.94-5).

McGinn, who holds a similar view of experience, suggests that,

[although] different sense-modalities may present the same kinds of environmental feature, e.g. shape or texture – as with sight and touch – ... the subjectively distinct experiences that present these features also present other features...[the] differences in the range of contents available to different types of experience seem enough to capture the obvious phenomenological differences in the experiences associated with different senses (1988, p.35).

Both McGinn and Dretske think that we can explain what makes the distinct experiences associated with each of the senses the kinds of experience they are by appealing to the kinds of properties that different experiences represent.

Whilst none of these writers is explicitly concerned with giving an account of the five senses, all of them apparently agree that there are distinct experiences of different kinds associated with each of the senses. They attempt to accommodate such differences within their general account of the nature of perceptual experience, and there is no obvious reason for thinking that they cannot in principle do so.

2. There does not, then, appear to be anything especially problematic about giving an account of the senses, if all such an account is an account of the different kinds of experience associated with each of the senses. What does seem problematic, however, is explaining the relation between different senses or, since we are thinking of the senses as constituted by different kinds of experience, the relation between the different kinds of experience associated with each of the

senses. We can illustrate why this is thought to be problematic by considering Molyneux's question.

Someone who thinks that there are distinct kinds of experience associated with each of the senses is likely to see a particular interpretation of Molyneux's question as central to questions about the senses, and one's answer to it as indicative of one's view of the nature of those experiences. In 1688 William Molyneux wrote to Locke (Locke 1689, II, ix, 35) and asked him whether a man who was born blind and who had learnt to distinguish shapes on the basis of touch could, if his sight were restored, say which shape was which by vision alone, without having touched them. Molyneux thought not, and Locke was inclined to agree with him. Historically, this question has been much discussed and many different issues have been seen, by different theorists, as central to its resolution.¹⁰

More recent discussions of Molyneux's question have understood it to be a question about the relation between sight and touch, about how we know that what we see is the same as what we touch. That is, as a question about our ability to recognise or classify instances of visually and tactually perceived shapes as the same kind of shape. One's answer to Molyneux's question can therefore be taken to reveal something about one's view of the character of the different kinds of experience associated with each of those senses, and the relation between them. Someone who answers No to the question will be someone who thinks that we must *learn* to relate the experiences of sight and touch (and of the different senses generally). If such a connection has to be learnt, that is because, as a consequence of the way that tactual experiences are different from visual experiences, shapes do not look the same way that they feel and so cannot simply be perceived to be the same. On the other hand, someone who answers Yes thinks that the experiences are such that we don't have to learn to relate them: things that have the same shape appear to us as having the same shape to both sight and touch, and so we can simply perceive them to have the same shape.

¹⁰ Locke and Molyneux saw it as a question about form perception and only tangentially a question about the senses. See Brandt Bolton (1994). For Berkeley the issue turned on abstractionism and the relation between the ideas of different senses. For a discussion, see Atherton (1990, Chs.5 and 10). Others see it as a question about innateness. For historical surveys and discussion see Degenaar (1996) and Morgan (1977).

Correspondingly, if one wants to decide which is the correct answer one needs to decide whether or not whatever it is that determines visual and tactual experiences as of distinct kinds is such that one must *learn* to relate the visual and tactual appearances of shapes. To answer Molyneux's question, then, the philosopher of perception needs to provide an account of our tactual awareness, and our visual awareness, of shape or space (or whatever other property we can perceive with two or more senses). She then must explain whether and how we can tell, on the basis of that awareness, that it is the same property that we are perceiving through sight and touch; that is, whether it is perceptually apparent to us that we are aware of the same property through sight and touch.

That is more or less Gareth Evans' approach to Molyneux's question; he takes "the issue on which Locke, Berkeley, Leibniz, Condillac, and others were taking up positions" to be "that of the relation between the perceptual representations of space attributable to the blind [i.e. via touch], and the perceptual representation of space available in visual perception" (1985a, p.370). He goes on to give an account of our tactual and visual experience of space which would imply a positive answer to Molyneux's question (pp.390 ff.). In a more recent discussion John Campbell (1996) presents the problem as one of explaining what he calls the "cross-modal transfer" of shape recognition – the ability we have to tell that a shape that we see is of the same kind as a shape that we touch. In particular he wants to know whether cross-modal transfer is a "rational phenomenon," that is, whether "the sameness of the property perceived in sight and touch is transparent to the subject" (p.304); that will depend, he thinks, on "whether there is a difference between the phenomenal characters of shape experience in sight and in touch" (p.301). He then argues for a particular account of shape perception which entails that there is no cross-modal difference in the phenomenal character of shape experience; and hence that cross-modal transfer is rational. It would follow that we should give a positive answer to Molyneux's question.

Dretske, in the passage quoted above, supposes that the fact that it is the same property that is represented by both visual and tactual experience is sufficient to explain how we are able to tell that we see and feel the same property: we perceive them as being the same. That is not inconsistent, he thinks, with the

experiences being of distinct kinds. Berkeley, of course, answers No to Molyneux's question because he holds that:

the extension, figures and notions perceived by sight are specifically distinct from the ideas of touch, called by the same names, nor is there any such thing as one idea, or kind of idea, common to both senses (1732, sec.127).

The ability to relate the ideas specific to one sense with those specific to another must be learnt. Some sense-datum theorists hold a similar view. Russell, for example, says that

space as we see it is not the same as space as we get it by the sense of touch; it is only by experience in infancy that we learn how to touch things we see, or how to get a sight of things which we feel touching us (1912, p.14).

Not all sense-datum theorists hold this kind of view. Robinson thinks that the question of whether properties like shape "appear in different senses in a way that is qualitatively exactly similar, or only similar in some more structural manner" is one to which there is no obvious answer (1994, p.207; see pp.207-11).

All of these writers are addressing the same problem. Given that there are different kinds of experience associated with each of the senses, how is it that we can tell that a property that we feel is the same as one that we see? They provide different answers to that question; answers that are usually determined by their view of the nature of perceptual experience generally. Given that there are these different kinds of experience that is a perfectly reasonable question to ask. But there is another question that we can ask. A question which is, it seems to me, more pressing.

It is rare that we perceive things with a single sense; it is far more usual for us to perceive things using two or more different senses simultaneously. We can see and touch and hear and smell one particular thing (a cat sitting on one's lap,

perhaps; or an old fashioned radio sitting on one's mantelpiece).¹¹ When we perceive something in this way we are aware of various of its properties using different senses, and yet we are aware of these properties as properties of a single particular thing. I can tell on the basis of perceiving it, for example, that the cat on my lap is black, furry and warm. How am I able to do this given that my perception of the cat and its properties involves different kinds of experience? In particular, how are the distinct particular experiences, which are instances of the different kinds of experience associated with each of the senses, related to one another such that I am aware of a single particular thing – the cat? The question which forms the focus of the discussions of Molyneux's question is how can experiences of different kinds provide me with an awareness that two things that I can perceive are of the same kind. The question I am asking here is how can two distinct particular experiences, instances of different kinds of experience, provide me with an awareness that what I perceive with each of them is a single particular thing?¹²

Berkeley and the sense-datum theorists have a straightforward answer to this question. According to them we are never, strictly speaking, aware of the same thing with two different senses. For Berkeley, our experience is constituted by an awareness of instances of sensible ideas, and the ideas of different senses are specifically distinct, so distinct senses can only ever provide us with an awareness of distinct instances of sensory ideas. Two experiences are distinct if and only if they are constituted by the awareness of distinct ideas; the experiences of different senses are distinct; so the experiences of different senses are constituted by the awareness of distinct ideas and we are never aware of the same particular thing with two senses. Something similar is true of the sense-datum theory. Experiences are constituted by an awareness of sense-data, and two experiences are distinct if and only if they are constituted by an awareness of distinct sense-data. If two particular experiences are distinct, they must be experiences of distinct particular sense-data. So if the experiences of different senses are distinct experiences they are experiences of distinct things, and we are never aware of the same thing with two or more

¹¹ Such radios give off quite a distinctive smell when they are switched on.

senses. Of course, both these views may go on to explain how it can seem to us that we are aware of the same thing; such appearances are, though, misleading. Although I don't have space to discuss whether and how they might do so,¹³ someone who holds a representational theory of perception would need, in similar vein, to explain how two distinct experiences can represent the same particular thing and represent it as being the same particular thing.

It is not clear that these solutions to the problem are satisfactory, but I'm not going to pursue these issues any further, partly because it is not clear what constraints there are on an answer to them, and so it is not clear what would constitute a satisfactory answer. It is not clear, for example, whether we should regard the sense-datum theorist's explanation of our awareness of particular objects as satisfactory, and if not why not. But mostly I'm not going to pursue them because I want to question the assumptions that got us this far; namely, that intuitively plausible view that there are distinct kinds of experience associated with each of the senses, and that an account of the senses must be an account of differences between those distinct kinds of experience. Further reflection on the character of our experience will, I suggest, undermine the intuitive plausibility of that assumption. A different problem then emerges for someone who wants to give an account of the five senses.

¹² Ayers sees a similar problem here, generated by what he calls "the fragmentation of the sense field". He sees this fragmentation as inevitable given a certain conception of the nature of our perceptual experience (Ayers 1993, Vol.1, pp.188 ff.).

¹³ But see Millikan (1997 and 1991).

2. The Unity of Experience

In the last chapter I suggested that there are, intuitively, distinct experiences of different kinds associated with each of the senses. If that were true, it would follow that we can explain what constitutes the different senses by appealing to the different kinds of experience involved in perceiving: one sees something in virtue of the kind of experience – a visual experience – involved in perceiving it, and so on for the other senses. I suggested too, that if one has this conception of the nature of the senses then the difficult problem in giving an account of them would seem to be that of explaining how they are integrated with one another. This conception of the nature of the senses generates an epistemological problem, that of telling the relation between what is perceived with different senses; but that there is this problem can seem compelling only given the background assumption that a correct account of the nature of the senses is in terms of the kinds of experience involved. If we don't make that assumption, then the fact that we have different senses is not something which makes this kind of recognition problematic; the question of how we recognise a property as the same when perceived with different senses will not, in that case, be a problem specific to the *senses*. Rather, it will be an instance of a more general problem about how we recognise the sameness of perceived properties over time and in different contexts.¹⁴

Are we right to make this assumption? In what follows I shall argue that we lack introspective support for the claim that there are distinct experiences of different kinds associated with each of the senses. When we reflect on the character of our perceptual experience it appears to be integrated across different senses. That means that, *prima facie*, there is no problem about telling whether what we perceive with one sense is the same or different to what we perceive with another; instead it generates a different kind of puzzle: given that our perceptual experience seems to us to be integrated, why and how do we distinguish different senses?¹⁵

¹⁴ See, for example, Millikan (1997). She discusses a general problem of what is involved in representing two things as identical, it is not a problem specific to the different senses.

¹⁵ Note that a problem about integration is generated by any 'atomistic' conception of experience which has to explain how simple ideas or perceptions get compounded into complex ideas of

Introspectively at least, there seems no basis for claiming that there are distinct experiences associated with each of the senses: the problem of the senses is, therefore, not how they are integrated with one another, but why and how we distinguish different senses at all.

1. What are the apparent objects, qualities, relations and so on, of each of the senses, and to what can we attend in attending to the objects of each of the senses? I suggest that, no matter with which sense the experience is produced, all of our sensory experience is perceptual and can provide us with an apparent awareness of objects (and the rest) existing independently of us.

Many writers think this is obviously true of vision. Harman, for example, describes a typical subject, Eloise, and claims that:

When Eloise sees a tree before her, the colours she experiences are all experienced as features of the tree and its surroundings. None of them are experienced as intrinsic features of her experience. Nor does she experience any features of anything as intrinsic features of her experience...There is nothing special about Eloise's visual experience. When you see a tree, you do not experience any features as intrinsic features of your experience. Look at a tree and try to turn your attention to the intrinsic features of your visual experience. I predict that the only features there to turn your attention to will be features of the presented tree...¹⁶

material objects. Again, this is not necessarily a problem about the senses. It isn't necessarily the case that there is a special problem in explaining how the simple ideas of one sense are combined with those of another which is different from the problem of explaining in general how simple ideas are combined. It is for this reason that we should see the problem Locke thinks addressed by Molyneux's question to be very different to the problem Berkeley sees as addressed by the same question. See Ayers (1993, Vol.1, Chs. 4 and 21); and see Brandt Bolton (1994) for Locke on Molyneux's question.

¹⁶ Harman (1990, p.39). For a similar expression of the same view, see Tye (1992, p.160).

In enjoying a visual experience we are apparently aware of objects and their properties as part of the world existing independent of us. How might we convince someone who rejected this claim?¹⁷

To claim that our visual experience provides us with an awareness of apparently independently existing objects is to claim that the things that we are aware of in having that experience appear to us to be independent of the experience in virtue of which we are aware of them. This is a claim about the phenomenology of experience, about how things seem or appear to us to be; we might attempt to justify it, therefore, by appealing to some apparent feature or property of the objects of our awareness which those objects couldn't have were they not independent of the experience. If the objects of our awareness have such features then they would be independent of our experience of them; if they appear to have such features, then they will appear to be independent of our experience.

Recently there have been several psychological studies of a visual phenomenon known as 'amodal completion'. These studies attempt to uncover the perceptual mechanisms underlying our perception of surfaces and make use of fairly sophisticated phenomenological tests. I think that we can see the results of these tests as lending support to the claim that we are apparently aware of independently existing objects and properties.¹⁸ In one such study, Nakayama et al (1995) showed subjects a diagram (p.9, figure 1.7. Reproduced below) of two apparently overlapping surfaces. Nearly everyone sees this figure in the same way – they see the surfaces y and z as connected and as passing behind x, and they see the border B as belonging to x and not to z. A natural description of this picture is that one surface appears to pass behind another. If that is right, then one surface appears to be obscured by another and so appears to be independent of our awareness of it.

¹⁷ I have in mind someone like Ayer (1973, p.91) who speaks of us being aware of "a visual leaf-pattern, a visual cat-pattern and so forth".




Figure 1.7

In another experiment subjects were shown a picture in which the several depictions of the letter B had been fragmented, in such a way that it is difficult to recognise them and difficult to tell how many instances of the letter are depicted (p.11, figure 1.8). When an occluder is added to the picture, so that the same fragments appear to be obscured, the letters is easily recognisable and can easily be counted. A plausible explanation of this difference in our ability to recognise the fragments is that with the occluder in place we see the letters as a complete surfaces – as several letter Bs – obscured by, and passing behind the occluder, whereas, without the occluder, we see only separate and unconnected fragments and not complete surfaces. What seems undeniable is that we see what appear to be a complete letters in the second figure. If we see what appear to be complete letters, then the surfaces of the letters must appear to us to pass out of sight behind the occluder and hence to exist independently of our awareness of them. This phenomenon is even more striking when differences in apparent depth are used to place the occluder either in front of or behind the plane of the fragmented letters (*ibid.*, p.12, figure 1.9). When the fragments appear to be behind the occluder we seem to see them as whole letters passing behind and obscured by it. When the fragments appear to be in front of the occluder they do not form complete surfaces and they are, as a consequence, difficult to recognise as fragments of the letter B. The same shape fragments look different to us in the two cases; and they do so

¹⁸ I should stress that this is not a question that the experimenters are asking, and so not a conclusion that they draw.

because in one case they appear to be complete surfaces passing behind something – and hence as existing unperceived. We can capture the difference in the way the letter fragments look by saying that in one case the letters appear to pass out of view behind the occluder. If that is a correct description of how the surfaces appear then our experience is of surfaces which appear to be independent of us.

It must be admitted that this argument is not conclusive. It may well be possible to explain the difference between those pictures in which the letters are easily recognisable and those in which they are not in some way which doesn't commit us to describing them as the appearance of independent surfaces. It might be said, therefore, that my description constitutes an over-description, and goes beyond what is strictly licensed by how things appear.

In the end, in answering this kind of objection, there is no better strategy than that employed by Strawson (1979). To describe what one's visual experience is like, he suggests, one should simply describe or characterise how things seem to one visually – what one's visual experience of things is like – at any time. If one attempts to faithfully describe the objects of one's visual experience in a way that does not distort or misrepresent their character, then one must employ concepts of independently existing objects and places. Sitting here, for example, it looks to me as though there is a desk in front of me with books and papers spread out over it and, through the window, there appears to be a garden with a clothes line running across it. Further away in the distance there look to be several trees, the nearer ones obscuring those that are more distant, all framed by a cloudy sky. It would misrepresent the character of my experience of these things if I were not to use these concepts of trees, books, and so on, in describing them. Strawson puts the point in this way:

Our perceptual judgements...embody or reflect a certain view of the world, as containing objects, variously propertied, located in a common space and continuing in their existence independently of our interrupted and relatively fleeting perceptions of them...it appears that we cannot give a veridical characterisation even of the sensible experience which these judgements...'go beyond',

without reference to those judgements themselves (1979, pp.94-6).

It would simply misrepresent the character of our visual experience to claim that we are not visually aware of apparent aspects of the world – objects, their qualities, relations, and so on.

Of course very few people deny this of vision. Even Ayer claims that we are only “implicitly” aware of leaf-patterns, cat patterns, and so forth (1973, p.91), and the proponents of sense-datum accounts of perception are perfectly aware of and would even endorse the kind of phenomenological observation made by Harman; what they dispute is the conclusion he draws from it. The same is not necessarily true of touch: the suggestion that we are tactually aware of apparently independently existing things can seem more open to dispute. One reason for this, I suspect, is the role of sensation in touch.

Why would anyone deny that in tactual perception we are aware of what appear to be independently existing objects?¹⁹ In his discussion of touch, Thomas Reid noted that there are two elements in tactual perception to which the subject can attend: a subjective sensation and an objective perception of the features of the object felt (1983, pp.35 ff.). He was surely right that sensation is involved in touch. When we touch an object a part of our body comes into contact with it and we can feel a sensation of contact; and we can attend to the sensation that we feel in that part of our body in contact with what we touch. Reid describes such sensations as subjective and internal to the mind, but that seems wrong: the sensation itself feels to have a location; it feels to be located in that part of the body that is in contact with what one touches.²⁰ That means that Reid’s account of the role of sensation in touch, as a subjective precursor or accompaniment to tactual perception, is mistaken. What role do bodily sensations play in tactual perception? One plausible answer to that question is that we are tactually aware of objects through being aware of the sensations in our body as it comes into contact with those objects:

¹⁹ Mill held such a view of tactual experience and perception, see (1979, pp.214 ff.).

²⁰ See Armstrong (1962) and Martin (1995) for a discussion and defence of the claim that bodily sensations are perceptions of the body.

The model of touch here is that of the body as template. We are embodied in a world which contains potentially many other bodies. We can come into contact with other bodies, and they can impede our movement and distort our shape. Such physical impingement on us is reflected in the awareness we have of our bodies. One is aware when one's movement is impeded, and when one's skin is in contact with objects or is distended by them. In being aware of one's body, sensing how it is disposed, where it can and can't move, and where one has sensations, one can attend to the objects in virtue of which these are true. One measures the properties of objects in the world around one against one's body. So in having an awareness of one's body one has a sense of touch (Martin 1992, p.203).

This account of the role of sensation in touch is quite consistent with the claim that in tactual perception we are apparently aware of independently existing objects and their features, so the fact that sensations are involved in touch gives us no reason for rejecting the claim that tactual experience provides us with an awareness of what appear to be independently existing objects. Indeed, if we attempt to characterise our tactual perception we will do so in terms of objects and their properties located in space around us.

Think of all the different things that one can perceive by using the sense of touch.²¹ To pick an everyday example, I can focus my attention on a cup of coffee I'm holding in my hand. Although my hand is only in contact with a part of the cup – the points at which my fingers touch it – I can nonetheless feel its shape and size and its orientation. The cup seems to extend beyond the points of contact with my fingers, but not indefinitely: it feels to have a definite bulk and extent, and I can tell that it is a thick, heavy, mug. That the cup itself feels to be of a definite size is probably due to the distribution of its weight and the way its centre of gravity shifts as I lift it (think of the difference between how long a pencil feels held between two

²¹ For descriptions of tactual perception, and its variety and richness, see Gibson (1966, chapters VI and VII), Krueger (1989), Lederman and Klatsky (1987), and Millar (1997).

fingers compared, say, to a length of wood), but perhaps has some visual input too (Aglioti, Goodale, and DeSouza 1995). I can feel that it is half full, just as I will be able to feel when it is empty, and I can feel that the coffee it contains is hot. There are various things that I can tell about the material from which the cup is made. I can feel the texture of its surface and that it is solid – a smooth ceramic – very different to the feel of a polystyrene cup, for example. Of course I can feel where the cup is and so I know exactly what to do – where to move it – in order to drink its contents (something that I can do with my eyes closed). Normally, of course, we don't pay very much attention to how a cup we are drinking from feels, but when we do it seems that, just as in the case of vision, a veridical characterisation of this awareness must characterise it as an apparent awareness of an object and its properties which appear to be independent of that awareness. My tactual awareness is not limited at any time to an awareness of what I am holding in my hand. As I sit at my desk I can feel the chair that I am sitting on – its shape, location, and so on – and the desk under my elbows. If I attend to them, I can feel (to a certain extent anyway) the clothes that I am wearing, especially when I move. I can feel the glasses that are resting on my nose. I can feel the floor under my feet. At any time my tactual awareness is of a world of things as they press up against me. The things of which I am aware appear to be just that: independently existing objects arranged in space around me and coming into contact with my body.

What kinds of thing do we hear? We talk of hearing people speak, of hearing aeroplanes flying overhead, of hearing dogs barking, and doors squeaking. We talk, that is, of hearing independently existing objects; and so it seems that we are auditorily aware of what are apparently independently existing things. But, it might be objected, we don't really appear to hear people, or aeroplanes, or doors; what we appear to hear is the sounds that they make. Even if that is true (and I'm not sure that it is) the sounds that we hear themselves appear to exist independently of our hearing them. I discuss sounds in detail in the next chapter, but the reasons for thinking this are, briefly, as follows. Sounds have apparent spatial properties: they can appear to come from certain locations in space, and to move; they are perhaps best thought of as a certain kind of object since, like objects, they are things that we can re-encounter over time. When, for example, we hear a shout echo back to us from a cliff or in a cave we can appear to hear again the very sound we heard

earlier. It's the very same sound that we re-encounter. If, therefore, one were to characterise what one hears one would describe sounds as things which appear to us as existing independently in the world.

2. I have argued, then, that a veridical characterisation of our visual, tactual, and auditory experience would characterise it as the experience of what appear to us as aspects of an independent spatial world: in having these different experiences we can become aware of things which appear to be independent of our experience of them. If true that is sufficient to generate the problem about the senses that I describe below, whether or not taste and smell can be similarly characterised.

Taste and smell are the two senses for which the claim that they provide an awareness of what are apparently independently existing things might seem least plausible. They are, it might be thought, the most subjective, least perceptual, of the five senses. Unlike the other senses which provide us with an awareness of things as being some way or other, when we experience a smell or enjoy a taste we are not aware of anything – there is nothing which appears to us to be some way – we are simply affected in a certain kind of way.

It would be wrong, however, to characterise the experience of taste as simply a matter of being affected in a certain kind of way. We experience tastes as located in our mouths, often as on the tongue. So even if there are aspects of our experience of taste which appear subjective, taste experiences are experiences of the tongue as being affected in a certain way.²²

On some occasions a taste or flavour can appear to be the flavour of something that we place in the mouth. In fact, what we think of as the flavour of

²² Sensations appear as located in parts of the body, so in having them we have experiences of certain parts of the body. But there may be some aspects of some sensations – sensations like, for example, pain – which are experienced as subjective. There may, that is, be aspects of what some sensations are like which are simply a matter of being affected in a certain way and which are not experienced as aspects of whatever part of the body we experience in having the sensation. For a discussion, see Martin (1998a) and see Ayers's discussion of pain as a secondary quality (1993, Vol 1, pp214-216); see also Peacocke on 'hurt' (1985). McGinn claims that "sensations do not have intentional objects in the way that perceptual experiences do" (1997, p.8). That's wrong: sensations do have intentional

something is frequently an amalgam of sensory experience in which smelling is an important ingredient. The taste of an apple is different to the taste of a piece of raw potato and we can easily distinguish them. It is much harder to distinguish them, however, if one's nose is blocked. Using the tongue alone they both just taste faintly sweet. There is also a strong tactile element involved in what we think of as taste, in the strong hot flavour of English mustard, for example, and in the texture that food feels to have in the mouth and when chewed. Even the sound that some foods make can be important:

Crisp foods have to be loud in the upper register. They have to produce a high-frequency shattering; foods which generate low-frequency rumblings are crunchy or slurpy but not crisp...²³

All these properties go to make up what we regard as the taste of something. The taste of something is a rich and complex experience of an amalgam of different properties. Someone who attempts to describe how something tastes may mention various things about the apparent flavour, texture, consistency, temperature, and so on, of the food they are tasting; whether it is hard or soft, whether it melts on their tongue or is chewy. They will say, too, whether it is sweet or bitter, and how it affects their mouth, whether it is spicy and produces a burning sensation, for example. This kind of description is a description of the properties of apparently independently existing things. A veridical characterisation of what it's like to taste something, therefore, would describe features of apparently independently existing objects.²⁴

objects – parts of the body; rejecting that claim is consistent with allowing that some aspects of our sensations are experienced as subjective and not as features of their intentional objects.

²³ Bodanis (1986) quoted in Ackerman (1990, p.142). See chapter 3 of Ackerman for an interesting description of the many different features of food that go to make up the way that it tastes. See the papers in Carterette and Friedman (1978) for a more detailed and rigorous discussion of taste perception.

²⁴ Someone might say: Strictly speaking, the experience of taste is just the experience of salt, sour, bitter, and sweet in different combinations. Such a claim would, however, be revisionary of our

Even the sense of smell provides us with an awareness of something: smells, or odours. Is it right, though, to say that we are aware *of* smells rather than that, in experiencing a smell, one is simply affected in a certain way?²⁵ Nothing in what follows will depend on an answer to that question, but I do think that there are reasons for thinking that our experience of smells is the experience of what appear to be independently existing, if somewhat impermanent and vaguely delineated, things. One can perceive a smell as occupying a certain place in the world (like a cloud perhaps) and as given off by something; a smell can appear to be in a jar of coffee, or spread over a field of poppies, or as contained in a room. One's experience of a smell can become more or less intense as one moves around or through it, and one can block one's experience of it by holding one's nose. Some of the smells we experience have properties which affect us in other ways – they might sting our nose and make us sneeze, for example. When they do so, we experience the stinging feeling as located in our nose, and the smell as something that causes that feeling.

The apparent objects of all of our senses, then, are independently existing things: aspects of the independent spatial world. In attending to the objects of our experience and their properties we are attending to things which exist independently of our experience of them. This is certainly true of vision, touch, and hearing; it is plausible too, I have suggested, of taste and smell. This is a claim about the phenomenology of our experience, about how things seem to us to be.

I have been describing the apparent objects of each of the five senses, considering each in isolation from the others. But again, it's rare that we perceive the world with a single sensory modality; usually we perceive the world with more than one sense simultaneously. Just as we can consider the objects of the experiences of each of the senses considered in isolation, so we can consider the objects of our experience of the senses when they are used in conjunction with one

common sense classifications of the senses. What we judge to be the taste of something includes much more than just the experiences of these properties.

²⁵ Lycan (1996, pp.144-51) argues that our olfactory experience is the experience of something objective: that it represents things in the world as being a certain way. For an alternative view of smell see Perkins (1983, Ch.3). For an introduction to the psychology of smell, see Gibson (1966, pp.144 ff.), and see Carterette and Friedman (1978, part III).

another. In particular, we can ask how things appear to us when we perceive them with more than one sense.

Suppose that this lunchtime you strolled from the philosophy department at University College to the Bloomsbury cafeteria to buy a cup of coffee: out of the front door of the department, along the street and then through the double doors and into the Bloomsbury. As you did this you may not have been paying much attention to what was going on around you. You would, nonetheless, have been aware of many things: we are continually and unreflectively aware of many aspects of our environment as we move around it. Imagine what it would be like. When you first stepped out the door the air would feel cold as it hit you in the face. You would immediately be engulfed by the sounds of the street, particularly the traffic as it runs around the edge of the square, and of people talking as they walk past. You would, of course, see these things too – a motorbike which you hear approach and speed past, the people whose voices you can hear, the rest of the street and the things in it, and other streets further away. As you walk down the steps and along the street you would feel the hardness of the ground beneath your feet (a completely different feel to that of, say, grass, or sand). As you walk down the street you would see and hear people approach and pass by; if it's busy you might brush into some of them. For a moment you would feel the cold glass against you hand as you push through the double doors into the Bloomsbury, then you would feel the warm air and smell the cigarette smoke; and the noise changes – there would be less traffic noise, more people talking, and that aural ambience distinctive of an enclosed space, and so on. (We could add detail to this description almost indefinitely.)

Notice that this awareness involved all of your several senses; by that I don't just mean that all your senses were working, but rather that you would have been aware of particular things with several of your senses. Take, for example, your awareness of the traffic: you can see, hear, and smell it; and perhaps you can feel it; or of the other people: you could hear and see them, maybe bump into them, perhaps smell them, too. Nearly all of our perceptual experience of things is multi-sensory, but we are not unreflectively aware of it as being that way. In being aware of the world we are not normally aware *of* our senses at all (they are not like windows out of which we peer), we are just aware of the things we perceive around us. This awareness of things which involves all the senses, is itself unitary in the

following sense: Whatever we are aware of we are – or can be – aware of as an aspect of a single spatial world. This is a feature of how things seem to us, of how we experience things as being.

Consider, too, your awareness of the apparent spatial location and relations between the things that you perceive. When you perceive something you often perceive its apparent spatial location – you can tell where it is. The objects and features whose apparent location we can perceive we perceive as having an apparent location within a single system of spatial relations. This is true of the objects and features we perceive with different senses. When we are aware of the apparent locations of things we are not, or not usually, aware of the objects we touch as apparently located in one space, of the objects we see as apparently located in another space, of the things we hear as apparently located in a third space, and so on, without being aware of the apparent spatial relations between them. The surface of the desk on which my arm is resting, for example, appears to be the same surface as that I can see; and the sound that I hear as I drum my fingers on this surface appears to come from where I can see and feel my fingers moving. Normally, in our everyday interactions with things, our attention is focussed on things in the world. When it is so, the locations of things we attend to through touch appear to be just those locations we attend to through sight and hearing.

That our experience of these spatial properties is as of them as unified or integrated across different senses is reflected in the abilities we have to act on the things that we perceive. Suppose you hear someone scream and you can tell where the scream seems to come from. You will know where to look if you want to see who screamed, and you know in which direction to point if you want to point at whoever screamed. When you look or point to a place from which a sound appears to come, it seems to you that you are looking at the same place you hear the sound to come from. In such cases our awareness is such that we don't have to work out how to move in order to point in the direction of something we can see or hear, or in order to reach to something we can feel.²⁶ There are, of course, exceptions to this. Sometimes we perceive something and, for one reason or another, we are unable to tell where it is. Our hand, to use an example of Peacocke's, may be

²⁶ A point emphasised by Evans (1982, pp.155 ff.), following Malpas (1968).

twisted around behind our back in such a way that although we can feel something with our fingers we may not be certain where it is that we feel it to be (c.f. Peacocke 1992, Ch.3), and we might be unsure, in such a case, where to move the other hand if we wanted to touch the first. Similarly, when someone is made to wear spectacles with inverting lenses, they have difficulty in co-ordinating the location of what they see with the other things they perceive. The very anomalousness of such cases merely emphasises how the objects of our experience are usually spatially integrated. (More significantly, these kinds of case do not occur for all and only things perceived using a single kind of sense modality, and so provide no reason for thinking that what we perceive with any one sense is not spatially integrated with the rest.)

Our perceptual experience is, then, as of a unified or integrated spatial world, and everything we perceive to have an apparent location is perceived to be apparently located in this space. The things that we perceive using each of our senses are such that we, for the most part, are aware of them as having an apparent location in a single space.

Might someone object that we don't experience things as integrated, rather our awareness of the relations between things is the result of judgements? The first thing that can be said in response to such an objection is that I am simply making a claim about the phenomenology of our experience, about how things seem to us. I am not giving any explanation of how or why they seem to us to be that way. It may be that there is an explanation to be given, and that such an explanation would appeal the fact that we make judgements; that would not show that my description of how things seem is wrong.

Actually, there is a reason to think that an explanation of the kind of sensory integration that I am describing will explain it as a feature of our experience, rather than a consequence of judgement. The reason is that there exist certain kinds of inter-sensory illusion.²⁷ For example, when someone's vision is displaced or inverted by their wearing prismatic glasses objects do not feel to be where they look to be (at least initially; if worn continually, adaptation occurs). Someone wearing

²⁷ Ayers (1993, Vol.1, pp.187-8) mentions some of these. For others see Welch (1978). In the next chapter I discuss a particular kind of inter-sensory interaction in much greater detail.

such glasses may be perfectly aware of the spatial relations amongst things that they can see, and of the spatial relations amongst things they perceive with their other senses, but they will be unable to relate the two. The systems of spatial relations between what they see and what they touch become dissociated (for a description of prism adaptation experiments see Gregory 1966, pp.204-14 and Rock 1975, pp.460 ff.). One result of this is that they don't know where to reach in order to grasp something that they can see. After a time, the experience of someone wearing such glasses changes and they adapt to them: sight and touch become re-integrated; even before they do so the subject might visually guide their hand and arm movements or work out how to move. What initially changes when someone puts on such glasses is their perceptual experience. They experience a kind of illusion: things look to be located in places that they are not in fact located. Someone who denied that the spatial integration of our senses is part of the content of our experience would have to say that the way things appear to someone wearing inverting glasses results from errors of judgement. But since the disintegration persists even when the subject knows that they are wearing glasses and are subject to an illusion, such cases display the characteristic feature of the content of experience as opposed to the content of judgement: they need not change when the subject has other information which leads to a judgement incompatible with content of the experience. So the fact that, when first wearing prismatic glasses, objects don't feel to be where they look to be, suggests that we that spatial integration is part of the content of our experience and not a consequence of judgments.

As I pointed out, in our everyday progress through the world our attention is normally drawn to the things – objects, places, and their features – that we perceive and whatever aspects of them engage our interests and intentions, we rarely attend to the various experiences we have of them. So, for example, I can perceive the desk at which I sit and the things in the room around me and my interest is in the objects themselves. My perception of them facilitates my everyday activities: I can move around the room, pick up a book and read it, write a note on a piece of paper, or drink from a cup of coffee. I can fairly easily describe the various objects that I can perceive: the various books that I see on the desk in front of me, the feel of the chair that I am sitting on, and the arrangement of flower pots in the garden outside the window. It takes more reflection and skill to go beyond such a

description and to describe the way these things appear to me, to describe my experience of these objects. To do so I must attend to and reflect on my experience in a particular kind of way.

When I look out of my window I can see the wall at the end of the yard. I can see the individual bricks that make it up, their colour and shape, and so on. As well as simply looking at the wall, I can reflect on my experience of looking at the wall and in doing so I can come to find out or notice various aspects of my experience.²⁸ When my attention is directed at the objects of my experience, the wall and yard occupy the focus of my attention; when I turn my attention inwards, to my experience, the wall is not replaced by any other objects or features which belong to a realm other than the realm of public objects of which the wall is a part. I discover what my experience of looking at the wall is like through perceptually attending to the wall and reflecting on that whilst I do it. In attending to my experience there is nothing for me to attend to but the objects of that experience. I find out about my experience by a certain kind of reflective attention to what it is an experience of. As Greg McCulloch remarks, “[r]eflection on the detailed character of visual experience is as world-directed an activity as ordinary seeing itself”; there are no “special...objects of introspection” (1993, p.53). What is true of visual experience is true of perceptual experience generally. In attending to our perceptual experience there is nothing for us to attend to other than the objects of that experience. This same point is made by Grice, when he says that

such experiences as...seeing and feeling seem to be, as it were, diaphanous: if we were asked to pay close attention, on a given occasion, to our seeing and feeling as distinct from what was being seen or felt, we should not know how to proceed; and the attempt to describe the differences between seeing and feeling seems to dissolve into a description of what we see and feel (1967, p.259).

²⁸ See Martin (1998b, Sec. II) for an extended discussion of attention to visual experience, in particular attention to the appearance of shadows.

Grice is telling us something about the character of our introspection of experience. When we introspect our experience all that we can attend to is the objects of that experience. I described the way that the objects of our experience are integrated with one another. In our everyday interactions with things, we are not aware of the things we perceive as occupying distinct worlds according to the sense with which we perceive them, we are aware of all the things we perceive as apparent objects and features of a single world. When we reflect on what our experience is like there is nothing to attend to except those objects and features. So, just as the objects of our experience are integrated, so when reflect on our experience we simply reflect on what it is like to perceive those objects and features. When we reflect on our perceptual experience in this way, it doesn't fragment into distinct experiences associated with different senses and we are not aware of having distinct experiences. There do not seem to us to be distinct experiences involved in perceiving something with more than one sense. If I am right about that, it undermines our intuition that there are distinct kinds of experience associated with each of the senses: when we introspect there don't appear to be distinct kinds of experience because there don't appear to be distinct experiences which could be grouped into kinds.

As well as undermining that intuition, it also creates a problem for the idea that we can distinguish the senses – i.e. explain what it is to see, say, rather than touch something – by appealing to distinct experiences of different kinds. It does so because we can not explain in virtue of what it is right to say that we are having distinct particular experiences when we perceive something with two or more senses. We can sometimes make sense of talking about two distinct and different experiences involved in perceiving with different senses, particularly in those kinds of cases in which the senses are used in isolation from one another.

We might, for example, contrast the kind of experience we have on a particular occasion of an object that we see but cannot feel with that we have of an object we can feel but cannot see. If asked to describe the character of our experience of these objects, to describe our visual experience in contrast to our tactual experience, we will describe different things. The features or aspects of an object that we are aware of when we see it are different to those of an object we are aware of when we feel it, so what we describe in describing the experiences will be different. That would be grounds for saying that our experiences of these distinct

objects are different. In describing the two experiences we will describe the two objects we perceive; since we perceive them with different senses, we will be aware of different features of these objects; the descriptions that we give of them will therefore be different. Alternatively, we might contrast what it's like to perceive things using only vision with what it is like to perceive things using only touch; that is, we might contrast the difference over time between the experience we have when we close our eyes and perceive things with touch only, and the experience we have when we avoid (as much as possible) using touch and simply look at things. When we do this, the way the world seems visually is different to the way the world seems tactually; if asked to describe how things seem to you in each case what you would describe would be different. The aspects of the world that you are aware of differs in each case, and so what you describe in describing your experience will be different. So again we can make sense of talking of two experiences being distinct (they occur at different times) and being different kinds of experience (they are of different kinds of thing).²⁹

But what sense can we make of our having or there being two distinct experiences in those cases in which we perceive the same thing with two or more senses simultaneously? On what basis could we claim that there are two distinct experiences involved when, for example, we are touching and looking at the same object? We can't contrast seeing one object with feeling another object since it appears to us to be the same object that we see and feel; if asked to describe what object we see and what object we feel we should describe the very same object. So we cannot say that we are having two distinct experiences because we are having experiences of two distinct objects. Nor can we contrast simply seeing the object with simply feeling it – we are doing both at the same time. So we cannot make sense of there being two distinct experiences in that way either.

In describing the feeling of the object it seems that you will simply describe these aspects of the object that you can feel; in describing the look of an object you

²⁹ This is what we have in mind when we think about what a blind person or deaf person lacks; we might say that, as a consequence of being blind or deaf, they lack the ability to enjoy a certain kind of experience of the world. But we don't have to describe them that way: they are simply unable to perceive certain aspects of the world.

will simply describe those aspects of the object that you can see. Rather than thinking of the senses as constituted by distinct and different experiences, it looks as though we should think of them as experiences of distinct aspects of the things we are aware of. *Prima facie*, we cannot, in the cases I am describing, explain the senses by appealing to different, distinct, experiences because we don't seem to be able to make sense of the idea of having distinct experiences of the objects that we perceive; we are simply aware of the object and its features. When we reflect on what this awareness is like, when we attempt to describe what our experience of the object is like, our description simple dissolves into a description of the object and its features. In attending to your visual awareness of an object you will attend to those aspects of it that you can see; when attending to your tactual awareness you will attend to those aspects that you can feel. You attend to different aspects of the same object. If asked to switch your attention between the experience of seeing and the experience of touching an object, all we can do is switch our attention between the things that we see and touch. At most, then, the distinct experiences associated with each of the senses are experiences of distinct aspects of what we perceive. An account of the distinction between the sense needs to explain, then, in virtue of what a perception of a particular thing is a perception involving a particular sense, and in virtue of what a perception of a particular property of something is, say, a visual perception of that property, and so on for the other senses. An account of the senses cannot simply be an account of what makes distinct experiences different, because it is not the case that every tactual perception of the shape of something involves a distinct experience to that involved in a visual perception of the shape of that thing.

The fact that our experience is unitary in the way that I described means that what we can appeal to in order to explain what makes experience visual when the senses are used independently of one another won't necessarily be sufficient to explain what makes our awareness of something visual when the senses are used together.

3. In the last chapter I described how many writers see the problem of the senses to be an integration problem – a problem about how we tell that what we perceive with one sense is the same as perceive with another sense. But when we reflect on

what our multi-sensory experience of the world and things in it is actually like, there is no such problem: the objects of our experience appear to be integrated. All the things that we are aware of we seem to be aware of as aspects of a single integrated spatial world. The problem of the senses is not an integration problem, it is rather a problem about how and why we distinguish different senses. What is it about our perceptions of some aspects of the world that lead us to distinguish them as seen or touched; what does, say, touching something consist in? That is the question that this thesis attempts to answer.

3. The Production of Sounds

In this chapter I describe our experience of a particular kind of thing – what I call the experience of the production of sounds. My aim in doing so is twofold. Firstly, I aim to provide a concrete example of the kind of unity of the senses that I described in the previous chapter. The experience of the production of sounds is an experience which essentially involves two senses; the fact that we have such experiences therefore supports my claim that our multi-sensory perception of the world does not comprise distinct experiences each of which is associated with one of the senses.

It follows that an explanation of the distinction we make between the senses cannot be an explanation of how such distinct experiences differ from one another. Secondly, I argue that the kind of experience that I describe generates a challenge for both the Berkeleian and sense-datum theorist's explanation of the distinction between the five senses. It doesn't directly challenge their accounts of the nature of experience, but questions whether those accounts can provide the kind of straightforward explanation of the distinction that I described. If, as I argue, they cannot do so, then adopting such an account of the nature of experience will be of no direct help to someone who wants to provide an explanation of the distinction that we make between the five senses. I spend most of this chapter characterising the kind of experience – of the production of sounds – that I am interested in. At the end I return to questions about the senses, and draw some specific conclusions.

In a recent book Jonathan Rée (1999) tells of how as a child he used to wonder which would be worse: to lose one's sight or one's hearing? Much worse, he concluded, to lose one's sight, since

Sounds seemed to me to be nature's waifs and strays: they did not fit into the familiar world of physical things, and they could not be tracked down by my other senses either...They were not part of the material world, and they had no weight to them, no substance. Is it surprising that I thought I could happily do without them? (p.19).

Whether or not we would be happy to do without sounds, the idea that our experience of sounds is of things which are distinct from the world of material objects can seem compelling. All you have to do to confirm it is close your eyes and reflect on the character of your auditory experience.

In what follows I will describe the features of our auditory experience which can lead one to think of sounds in this way. I will then describe a way in which we can experience sounds to be part of the material world; we experience sounds, I shall argue, as produced by material things.

1. Strawson, in his famous discussion of auditory experience, thinks it obvious that “where experience is supposed to be exclusively auditory in character, there would [not] be any place for spatial concepts...The only objects of sense-experience would be sounds. Sounds of course have temporal relations to each other, and may vary in character in certain ways...But they have no intrinsic spatial characteristics: such expressions as ‘to the left of’, ‘spatially above’, ‘nearer’, ‘farther’ have no intrinsic auditory significance.” (1959, p.65.).

We can, he thinks, adequately characterise our auditory experience in terms of auditory features alone and so purely auditory experience is the experience of what he calls a “No-Space world”. Since we think of the material world as a spatial world, if Strawson is right that our auditory experience is not intrinsically spatial then it is unsurprising that sounds can seem to be in a world apart from the material world rather than part of it.

Why does Strawson think it obvious that sounds have no intrinsic spatial properties and hence that there is no place for spatial concepts in auditory experience? Such a claim might strike one as simply mistaken, resting, one might think, on a rather thin conception of auditory experience or on a very thick conception of spatial properties. After all, one can hear sounds as coming from particular directions, or as occupying certain places; even with your eyes closed you can tell that my voice is coming from over here, and as a result you can hear where I am. Of course, one would have to admit that the auditory world is spatially impoverished relative to the visual world; sounds do not have the spatial dimensions of the material objects of vision. But it doesn't follow from that that the world of sounds is a No-Space world.

Strawson of course recognises that we can hear sounds as having a location, but our doing so is, he says, explained by “the existence of correlations between variations of which sound is intrinsically capable and other non-auditory features of our sense experience...the *de facto* existence of such correlations is a necessary condition of our assigning distances and directions as we do on the strength of hearing alone.” So even though we think and speak of sounds as having spatial properties, they do not have them intrinsically. But suppose someone were simply to deny this? What is it about auditory experience which warrants Strawson’s conclusion? He clearly thinks that there is a contrast between auditory experience and visual experience which *is* intrinsically spatial. What is this contrast with visual experience? The visual field, he says, “is necessarily extended at any moment, and its parts must exhibit spatial relations to each other” (p.65). So although a purely visual concept of space might be “impoverished compared with our own... it [is not] an impossibility. A purely auditory concept of space, on the other hand, is an impossibility” (pp.65-6).³⁰ There is, then, some contrast between visual experience, in particular the experience of a visual field, and auditory experience, from which it follows that visual experience is intrinsically spatial and auditory experience is not.

In saying that vision involves a visual field one needn’t be committed to any particular theory of perception; in particular one needn’t accept a sense datum theory according to which the visual field is constituted by some mind-dependent array of colour patches of which the subject is aware and through which she becomes aware of the world. The claim that vision involves a visual field can instead be understood to be indicating certain phenomenological features of visual experience, features which can be identified independently of any particular theory of perception. In “Sight and Touch” M.G.F. Martin provides a fine description of the relevant features:

³⁰ One might think that a visual concept of space would be impoverished because our concept of space, as Strawson suggests, is the idea “of a spatial system of objects, through which oneself, another object, moves, but which extends beyond the limits of one’s observation at any moment, or, more generally, is never fully revealed to observation at any moment”; it is a conception of something objective – capable of existing independent of one’s experience of it. A conception of space in this sense cannot be given purely visual significance; it is, rather, a theoretical conception (c.f. Evans 1985b).

we can think of normal visual experience as experience not only of objects which are located in some space, but as of a space within which they are located. The space is part of the experience in as much as one is aware of the region as a potential location for objects of vision. This is not to say that one can actually experience all sub-regions of a visual space at one time...objects occlude each other. The occluded areas of the visual scene count as part of the visual space in the sense that one could come to be aware of something at that location without altering the limits of the visual field provided by the angle of vision at that time. An area can come into view simply by a re-arrangement of things within the field, rather than by changing the field itself (1992, p.199).

He illustrates this with an example:

Consider the case of looking at...a polo mint...head on. One is aware of the various white parts of the mint arranged in a circle, and aware of how they are related to each another. One is also aware of the hole in the middle of the mint, and that that hole is there in the middle. If one was not aware of the hole one would not see the mint to be a ring-shape rather than a circle. Nothing need be perceived within the hole. One is aware of the hole as a place where something potentially could be seen, not as where something is actually seen to be (p.199).

Martin goes on to contrast this feature of visual experience with tactual experience;³¹ but since we are interested in the spatial features of sounds, we can ask whether our

³¹ He points out that this picture does not apply to many examples of touch. When one grasps the rim of a cup with one's fingers one can become aware of the shape of the cup, that it is circular, even though one comes into contact with it only at the points where one's fingertips touch it. Here we can draw a contrast between the way in which one is aware of the points of the rim between the points of contact, and the points of contact themselves. Does the space within the rim fall within a tactual field? Are they potential objects of tactual awareness? In the case of vision (and of hearing)

auditory experience is the experience of an auditory space in the way our visual experience is the experience of a visual space. There doesn't seem to be an obvious way to answer this question. There is nothing in our auditory experience analogous to the bounded field of vision: we are not auditorily aware of there being any boundary to the region of space within which we can potentially hear sounds in the way that we are aware of there being a boundary to visual space, and so we can't draw a contrast between sounds which fall within and sounds which fall outside the auditory field. But the absence of such a boundary would seem to provide us with no grounds for denying the intrinsic spatiality of sounds since we are, it might be said, auditorily aware of a region of space as a potential location of sounds: it's a region of space centred on the head and stretching out all around for an indeterminate distance. We can hear sounds as falling within this region of space, and we can be aware of many sounds simultaneously occupying different places within this space, and of the spatial relations between them. We are aware too that there are places where we cannot actually hear any sounds, but where we could potentially hear sounds. So we are auditorily aware of a region of space within which we could potentially hear sounds which is analogous to our visual awareness of a region of space as a potential location for objects of vision. There is an auditory space which is just like visual space, except that it is stretched so as to be boundless.

Whilst it is true that there is a space within which we can potentially hear sounds, there is nonetheless a difference between the way that one is visually aware of places where something potentially could be seen and the way that one is auditorily aware of places where something could potentially be heard. In fact, there's a sense in which one is not auditorily aware of places where something could be heard at all.

In the case of vision, we can distinguish between having an experience of nothing at a place where we could experience something, and not having an experience of something in a place we could experience something. In being aware of the hole in the polo mint one has an experience of a place at which there is

we might answer this by asking whether we could experience something there. See Martin (1992, p.200).

nothing. But equally there are places within the visual field one is simply unaware of because they are obscured and one may be unaware that there is nothing at that place, and yet that place still falls within one's visual field in virtue of the fact that "[a]n area can come into view simply by a re-arrangement of things within the field, rather than by changing the field itself". It is this visual awareness of places where there is nothing which has no auditory equivalent.

We are not auditorily aware of empty places – there's no difference between not having an experience of a sound at some place, and experiencing no sound there. One may hear nothing at some place, but one does not thereby hear the place as a place that, were something to be there, one would hear it. We can imagine a kind of spatial hearing deficit which prevents one from hearing anything at a particular place relative to oneself. In such a case it would no longer be true that one could potentially hear sounds from that place, and yet such diminution of the auditory field would not be apparent within one's auditory experience.

I think that Strawson may have something like this in mind when he claims that sounds are not intrinsically spatial. We simply do not hear the space within which we hear sounds; this contrasts with vision since we *can* see the space within which we see objects. Still, we might think that drawing this contrast with visual experience doesn't really show that sounds are not intrinsically spatial. There are, after all, other kinds of experience – tactual experience, for example – whose spatial structure contrasts with that of visual experience and yet which we think is, or at least can be, intrinsically spatial; it doesn't generally follow that an experience is non-spatial because it lacks the spatial structure of visual experience.

Such an objection would miss the point of the contrast. It's not merely that auditory experience has a different spatial structure, but that its structure is not intrinsically spatial. We cannot hear a place except by hearing a sound at that place and so the space in which we hear sounds to be is not itself given in auditory experience. We are visually aware of the space within which we see things in a way that we are not auditorily aware of the space in which we hear sounds: spatial relations are not themselves given in auditory experience.³²

³² In claiming on this basis that sounds are not intrinsically spatial I am saying nothing inconsistent with Evans account (1982) of the spatial content of auditory experience. Our auditory experience,

In denying the intrinsic spatiality of sounds we needn't simply appeal to the lack of an auditory field in auditory experience: there is a further feature of sounds which supports such a claim. Sounds need have no spatial properties at all; that is, one can hear a sound without hearing it to have any spatial properties. Not only do some sounds apparently lack spatial properties, but sounds which we hear to have spatial properties can lose them. One might hear a sound, a low rumbling say, which initially one hears to come from a certain direction but which, as one listens, one ceases to hear as coming from anywhere. The same thing can happen in reverse: a sound that initially appears not to come from anywhere can gain an apparent location. We shouldn't think of these cases as simply a matter of a sound's spatial properties becoming more or less determinate. Although sometimes sounds can surround or engulf one, so that they appear to have no determinate location – the sound of an audience applauding around you, for example – many sounds which have no apparent spatial properties do not sound like that.³³

This is reflected in the way we identify or individuate sounds. We do not individuate sounds spatially, which means that sounds can maintain their identity even when they lose their spatial properties. Suppose that outside in the street you hear a bulldog and a pekingese start to fight; you hear the whole thing from the initial growls and yaps to the final yowls, and can tell exactly when one dog leaves off and the other begins. In such a case you can hear the sounds made by both dogs, and you can selectively attend to either; but there seems nothing spatial about picking out and attending to one sound rather than the other: you may not be able to tell very much at all about where the two dogs are.³⁴ Similarly, you might single

on his account, has spatial content in virtue of its connections with spatially directed behaviour; in virtue, that is, of connections to something extrinsic to our auditory experience. You might think that Evans shows that we really do hear sounds as having locations, coming from certain directions, and so on: that's part of the content of our auditory experience. That's quite consistent with the features of our experience of sounds that I am describing. It's just one explanation of what it is for a sound to have spatial properties. I am not denying that they can have them. See Evans (1982, pp.156-7).

³³ You can imagine hearing a sound, a particular piano recital, for example, without imagining the sound to come from anywhere. The same is not true of visualising: when you imagine seeing something you imagine it from some – not necessarily your own – spatial point of view.

³⁴ The example is from Campbell (1997, pp. 65-66).

out and attend to the sound of the oboe when listening to a piece of orchestral music without in any way being able to distinguish its location from the location of the other instruments that you can hear. In both cases you can identify and selectively attend to a particular sound without being able to pick out its spatial properties.

Another nice illustration of the non-spatial way we identify sounds is provided by a phenomenon known as “the auditory stream effect”. To produce the effect, listeners are presented with an endlessly repeated sequence of six different tones, alternating three high ones and three low ones. When the tones are presented slowly to the listeners they hear the sequence in the order in which they occur. When the tones are presented at a faster rate the listeners no longer hear the tones in the correct order, but instead hear two streams of tones, one containing the repeating cycle of the three low pitched tones, the other the cycle of high pitched tones. The single sequence appears to break up into two different sounds,³⁵ and yet there is no apparent change in the spatial properties of the sounds; the sounds may appear not to have any spatial properties at all (Bregman 1990, pp.17-18).

We can experience sounds as located and as losing their spatial properties – as ceasing to be located, so sounds are only contingently spatial. That’s not merely to say that, like the material objects we perceive, they have the particular spatial properties they do contingently, but that they could lack spatial properties altogether. That’s not true of vision. Although we can see something move, and so see its spatial properties change, we cannot see it lose them altogether; so the objects of our visual experience and our visual experience of them are both essentially spatial. This is not true of sounds: we can hear a sound without hearing it to have any spatial properties at all. When we do experience a sound to have spatial properties, that it has them at that time appears equally contingent; we can easily suppose that we could have heard that very sound without hearing it to have the spatial properties we actually hear it as having: although in fact we hear it to be

³⁵ Subjects were unable to focus their attention on both streams at the same time: when they focussed on one, the other was heard as vague background; as a consequence, listeners were unable to report the order of the six tones taken together, although they could report the order of the high or low tones individually.

located, it might not have been.³⁶ Not only can we imagine a world of sounds which is a no-space world but, I suggest, we can imagine the actual world of sounds as a no-space world.

When we come to reflect on our auditory experience, on what sounds are like, then it can seem that their connection with the objects of sight and touch is a precarious one. Sounds may appear to be, as Strawson says, correlated with the material world, but they do not appear to be part of it. We can imagine a world of sounds which is dissociated from the world of material objects; we can imagine, too, the sounds we actually hear apart from the things that we see and touch. There appears to be nothing intrinsic to the sounds that we actually hear to connect them with the world of sight and touch.

Compare this with our experience of the objects of sight and touch. It's much harder to imagine the visual world – the world of seen objects – as dissociated from the tactual world. We are visually and tactually aware of intrinsically spatial objects that we perceive to occupy a single integrated spatial world, and so in imagining the world of sight dissociated from the world of touch we must imagine something quite different from our ordinary experience; there's no mere correlation between the objects we see and feel: they are one and the same; you can see the very same object that you feel, and you can perceive that it is the same object. When one sees and feels a particular object in this way it's not possible to imagine the very object one feels apart from the object one sees; there appears to be a unity to the objects of sight and touch. So whilst sounds appear not to be part of the material world, the same is not true of the objects of sight and touch.

³⁶ A question that we might ask is: is the non-spatiality of sounds a consequence of the nature of our experience of sounds or of the nature of sounds themselves? In the case of vision we think of both the objects and the experiences as being spatial: the objects because they form part of the furniture of a spatial world, the experience because it is able to provide us with an intrinsically spatial awareness of that world. Even non-spatial things like after-images are experienced as spatial. That's a consequence of the spatial structure of visual experience. It looks like it's necessary that something that is visually perceived is perceived as having spatial properties. That's not true of sounds. Is that because although our experience of sounds is not intrinsically spatial sounds are spatial? Or are sounds non-spatial?

2. So far I have been talking about what our auditory experience is like and I have been assuming, implicitly, that in attending to one's auditory experience one attends to that alone, and that one ignores or excludes as much as possible the experiences of the other senses, perhaps by shutting one's eyes. But is this the right way to go about things? Is what is true of auditory experience alone true of the experience we have of sounds when we use our hearing together with sight and touch? We tend to assume that it is.

In "The Guermites Way", Proust describes Swann as he sits quietly waiting for Saint-Loup, reflecting, as he does so, on the sounds that he can hear:

I heard the tick of Saint-Loup's watch, which could not be far away. This tick changed place every moment, for I could not see the watch; it seemed to come from behind, from in front of me, from my right, from my left, sometimes to die away as though it were a long way off. Suddenly I caught sight of the watch on the table. Then I heard the tick in a fixed place from which it did not move again. That is to say, I thought I heard it at this place; I did not hear it there, I saw it there, for sounds have no position in space (1925, p.72).

You might have had a similar experience; it's one of a number of perceptual phenomena which reflect the way in which vision can dominate auditory perception.³⁷ You will be familiar with another example: the ventriloquism effect. The ventriloquism effect occurs when a voice in one place – that of the ventriloquist – is heard to come from a different place – the mouth of a dummy. By keeping her own mouth shut and directing our attention to the dummy whose mouth, eyes, lips, and other movements are manipulated in synchrony with what is being said, the ventriloquist is able to take advantage of the interaction between the viewer's different senses to create the illusion that the dummy is speaking: it can seem as if it is the dummy that is producing the sounds.

³⁷ Similar 'intersensory bias' effects can occur with various combinations of senses; another example – from the laboratory – is visual capture, in which the felt position of one's finger or hand is strongly

Although there is a spatial discrepancy between the visual and auditory sources of the perceived event, one hears the sound to be located in the direction of the visual source, and one perceives no discrepancy between what one perceives with different senses. A similar effect can occur when listening to someone speak from some distance in a crowded or echoing room: when looking at them speak you may hear the voice as coming from the speaker – you both hear and see where the source of the sound is; but when your eyes are turned elsewhere you may no longer be able to tell where the voice is coming from. With your eyes again fixed on the speaker, and on the movement of her lips, a clear sense of the source of the sound will return. This effect can occur even when the sound alone has no apparent location as when, for example, the speaker's voice is relayed to speakers positioned all around the room.

The persistence of the ventriloquism effect has been experimentally tested (Warren et al 1981). In the experiment, subjects (whose heads were held in position) viewed a monitor on which was played a video of a talking head. A soundtrack was played through a separate loud-speaker whose position relative to the monitor could be varied. Subjects were asked to judge whether the voice they could hear was that of the talking head and to indicate the direction in which they perceived it or, if they perceived them to be different, the direction in which they could see the head and the direction from which they could hear the voice. The experimenters varied the setup so that sometimes the recorded voice matched the mouth movements of the head; at other times they introduced a discrepancy between the voice and the mouth movements either by means of a short time delay between the two or by playing recordings in which different words were spoken by the head on the video and played over the loud-speaker. The relative location of, and the distance between, the monitor and the speaker were also varied.

The (surprising) result of this experiment was that subjects often perceived a single event – a head speaking – when the voice and picture were perceived to correspond with one another even when the actual source of the voice was different to that of the head; and when asked to indicate the direction of the head and the

biased in the direction of its prismatically displaced visual image. For some psychological studies see Welch (1978) and Stein and Meredith (1993, Ch.1).

voice they indicated the same direction, but got both wrong. That is, when asked in which direction they heard the voice to come from they indicated a position between that of the actual position of the voice and the actual position of the head (though much nearer to the actual position of the head); and when asked to indicate the direction in which they saw the head they indicated the same position as they did when asked the direction of the sound. When shown the head or played the voice separately and asked to indicate their position, subjects were generally accurate.

The first thing to note about the ventriloquism effect is that it is a perceptual phenomenon, and not the result of a judgement to the effect that what one sees is what one hears. The interaction of vision and hearing produces a change in the one's experience; the place a sound appears to come from when one simply hears the sound is different to the place it appears to come from when one both hears the sound and sees its source. One experiences a kind of perceptual illusion: it seems as though the sound is at the same place as the thing one sees – the talking head – even when it is not. This is brought out nicely by the experiment. The subjects of that experiment, when asked what they could perceive, reported that they experienced a single event – a person talking – and they failed to report any discrepancy between what they perceived with different senses even when there was such a discrepancy. If these subjects were merely judging that they were perceiving a single event, and ignoring or setting aside certain discrepancies between what they perceived with each sense, then we would expect them to be aware, or capable of becoming aware, of the discrepancy. Yet when subjects were told of the discrepancy between their senses they were often unable correct their responses. And, in general, knowing that one is being tricked by a ventriloquist does not reduce the effect, you may be knowingly entertained by a ventriloquist's skill. Thus the phenomenon that I am describing displays, in Peacocke's words, "the characteristic feature of the content of experience as opposed to the content of judgement in that [it] need not alter when additional information results in a judgement of a content incompatible with that of the experience" (1986, p.156). This appears to be a feature of all kinds of intersensory bias: the subjects of biased experiences find it very hard to correct their mistakes.

If the phenomenon is perceptual then how should we characterise the resultant experience? How does the bi-modal experience differ from hearing alone, or from the kind of experience we have when we see a bad ventriloquist who is unable to produce the effect? Is there simply a change in the apparent location of what is perceived with each sense? And how, in general, does the experience of seeing someone speak differ from simply hearing them? Do we simply hear their voice as coming from the same place that we see them to be? It would be wrong to characterise the change in one's experience as simply a change in its spatial properties. There is a difference between, on one hand, hearing a sound as coming from the same place as an event one sees, and on the other, the kind of experience one has of the ventriloquism effect or, more generally, the experience one has of seeing someone speak.

When we hear a sound we are often in no doubt about what produced it; not just what kind of sound it is, or what kind of thing produced it, but which particular happening did; this is something that we can perceive. We often see something happen and hear a sound, and we perceive the sound to have been produced by what we saw happen. In the experiment that I described, subjects experience the voice they hear being produced by the head that they see, they experience the head as responsible for the voice they hear. Similarly, when one experiences the ventriloquism effect, the voice one hears appears to be produced by the dummy; it's not just that one hears the voice as coming from the same place as one sees the dummy to be, one experiences the dummy and its mouth movements as responsible for what one hears.

We are familiar with experiences of seeing one event cause another, as Peacocke says:

anyone who sees the child's hand knocking over the tower of blocks, or a fork-lift truck as lifting a crate, has [experiences as of one event causing another]. These experiences would not be adequately characterised as seeing an event of one type following an event of another type. Rather, taking the experiences at face value, one would be disposed to judge that the child's movement caused the tower to fall

over or to judge that rising of the fork-lift truck's arms caused the crate to go up (1986, p.156).

In the kinds of case Peacocke describes, apparent causation between two events is visually perceived; I am suggesting that we can have similar experiences as of causation between things perceived with different senses. When we see a dog bark and hear the sound it makes we don't just hear a sound as coming from the same place we see the dog barking; we perceive the dog to be producing the sound we hear. When we see a hammer striking an anvil and hear the sound of the blow we perceive the hammer blow as producing the sound. We need not restrict the examples to sight and hearing; we can feel something produce a sound we hear, too: when we feel a tuning fork vibrate and hear the noise it makes, we perceive the noise as being produced by the fork's vibration.

These are all cases of perceiving some event followed by a sound, and of perceiving the sound to be produced by, or caused by, the event. This is something that we perceive, it is part of the content of our experience; *these* experiences would not be adequately characterised as hearing a sound to come from the same place as one sees something happen. If one takes one's experience of hearing someone one sees speak at face value, one would be disposed to judge that the person speaking is responsible for the sounds that one hears, that they are *producing* or *causing* those sounds.³⁸

But is this really something we experience, rather than simply something that we are disposed to judge as being the case in the right circumstances? It might be objected that we can characterise the kinds of experience I am describing without mentioning causation or production; the experience that I describe as the experience of something causing or producing a sound is really, it might be said, nothing more than the experience of hearing a sound to be located in the same place and to occur at the same time as an event we see.

³⁸ For the claim with respect to vision, see Michotte (1963, esp. appendix 2); and see Bruce and Green (1990, p.333) for a discussion and other references. Bruce and Green are sceptical of Michotte's claim that causality is *directly* perceived, but not of the claim that we do have experiences of the sort described by Michotte. It's just that they think that an explanation of our experience's representing causality must appeal to computations or inferences performed by the visual system

There is, however, a difference between hearing something to occur in the same place and at the same time as something we see happen, and perceiving a sound as being produced by something we see happen. Someone who denies we have the kind of experience I am describing won't be able to mark this difference. One can often hear sounds at the same place as one sees something happening without it seeming to one that the sound is produced by what one sees happening. This is well illustrated by the – incredibly irritating – experience we have all had of the sound track of a film being slightly out of synchronisation with the pictures. Even when we hear the voice as coming from the same place as the person we see on the screen – when, for example, the soundtrack is in stereo – it doesn't appear to us that the person is speaking the words, that they are responsible for them; that's why the experience is so irritating. Although we hear the voice as coming from the same place as something we see happening, we don't hear it as being produced by what we see happening. Knowing that the words we hear are being spoken by the person we see doesn't reduce the effect: it's judgement independent, something we experience.

Consider again the ventriloquism effect. One's experience of that effect is susceptible to what the psychologists who performed the experiment I described call the "cognitive compellingness of the stimulus situation". In effect, the more the thing you see appears to be the source of what you hear the stronger the effect, and something you see appears to be the source of what you hear when they match; when, in other words, it looks as if it is producing the sound. But why would there need to be such a match if all that we experienced as a result of the effect was spatio-temporal coincidence? The best explanation of why match is important, I suggest, is that what we experience as a result is not simply spatio-temporal coincidence of source and sound, but the production of the sound by the source. It is easier to produce an illusion of the sound as produced by something seen when the sound matches what is seen.

Someone might still object to my description of these experiences as perceptions as of causation for the reason that we can never have such experiences. Our concept of causation, it might be argued, is the concept of a kind of relation which we could not simply perceive to be instantiated. Peter Menzies, for example, suggests that the counterfactuals involved in an instance of causation make it a

relation that 'cannot plausibly be claimed to be an object of direct awareness' on the grounds that the truth of a counterfactual cannot be perceived (1993, pp.202-3). The conception of a cause to which I am appealing, however, is not this very general philosophical concept; it is, rather, one of a whole range of causal concepts that feature in our everyday thought and language, concepts which include: scrape, push, carry, knock over, squash, make and so on.³⁹ For as long as we allow that people possess and use such concepts, and can apply them to things on the basis of perceiving the interactions between, then we should allow that causality, in this sense, can be perceived. And people do possess and use such concepts; Anscombe is right when she says

As surely as we learned to call people by name or to report from seeing it that the cat was on the table, we also learned to report from having observed it that someone drank up the milk or that the dog made a funny noise or that things were cut or broken by whatever we saw cut or break them (1971, p.69).

In claiming that we experience causation between events and sounds I am claiming no more than that we experience relations such as these.⁴⁰

There is a difference between hearing the sound produced by something without seeing (or perceiving in some other way) whatever it is that produces it, and having the experience of a sound as being produced by something that you can see (or otherwise perceive). This latter experience is one that we can only enjoy as a result of using two different senses. We never simply *hear* something as producing a sound because we can't hear the sources of sounds apart from hearing the sounds that they make, and we never simply *see* sounds as being produced because we don't simply see sounds, but we can perceive sounds as being produced when we both see

³⁹ This list is from Anscombe (1971, pp.68-9).

⁴⁰ Anscombe points out that the apparent perception of such things may only be apparent: we may be deceived by false appearances. It should be noted, too, that we can accept that we have such experiences of causation without committing ourselves to any particular account of the nature of causation in the world (c.f. Peacocke 1986, p.156). For further discussion of these issues, see Armstrong (1997, pp.211 ff.).

the source of a sound and hear the sound that it is producing. This is the case when we see people speak; but it's not restricted to the perception of speaking: we can experience the noise of a hammer blow as being produced by the hammer blow we see.

3. Someone might accept that we have experiences of the kinds that I am describing but claim that such experiences do not require the use of both hearing and vision: one can experience things as producing sounds by hearing alone. We often talk of hearing the source of the sounds that we hear, of hearing the thing which produces the sound. So, for example, we often say that we can hear the dog barking as well as the dog's bark; and this doesn't appear simply to be a manner of speaking since we can *think* demonstratively about the sources of the sounds we hear. Many writers have drawn this conclusion. Campbell, in a discussion of perceptual demonstratives, claims that we can refer "to an ordinary physical object" on the basis of hearing the sound that it produces. So when you hear a dog barking you can refer, using demonstratives such as "that bulldog", to source of the sound – the dog itself (1997, p.65). Similarly, Moreland Perkins thinks that one can be auditorily aware of a person, and that one's auditory awareness of that person is in part constituted by an awareness of the sound they make: "hearing [his wife] speak from the kitchen, [a man] not only hears the sound she makes, he also, as we say, hears *her*: his...awareness of his wife we conceive also to be auditory. But auditory awareness is also sensory awareness; so when he hears her voice his...awareness of his wife is sensory awareness." (1983, p.12). These are all examples of purely auditory awareness of the source of a sound. We are aware of the source by being aware of the sound that it makes, and only through being aware of that sound. It is the kind of awareness of the source of a sound that you can have even when you don't see or otherwise perceive it.

Given that one can hear material objects, then it is possible to explain the distinction between hearing a sound as produced by an object and hearing a sound as not produced by an object differently: hearing a sound as produced is hearing the sound and the object in the way in which awareness of an object is partly constituted by awareness of the sound. When one does not recognise the sound then one does not hear the object; when one does recognise it then one does hear

the object. So we don't hear all sounds as produced by objects, but some we do. Vision might have a role to play in helping us recognise sounds, but it doesn't have the role that I claim it to have.

I don't want to deny that we can hear objects as well as the sounds that they make, but I do deny that the experience of the production of a sound is purely auditory. There are therefore two responses that might be made to this kind of objection.

The first point to make is that we shouldn't think that we can hear the sources of the sounds we hear in virtue of the fact that sounds are properties or features of their sources. Several writers seem to endorse the view that sounds are sensible qualities of their sources, and sometimes we do talk as if they were: we talk, for example, of a floorboard creaking, or of a bell as having a mellow sound. Hearing a sound, according to this kind of view, can be hearing the source of a sound in the same way that seeing a shape is seeing the shape of something (of an object or surface).⁴¹

Sounds are not properties or features of their sources, and we shouldn't think of our awareness of the source of a sound as being analogous to our awareness of an object whose shape we see. Unlike our awareness of the shape of an object, it is possible to hear a sound without hearing its source. Sounds are, moreover, distinct from their sources: one of Newton's minor triumphs in the *Principia* was a derivation of the velocity of sound. To check his derivation he measured the time for an echo to return from the end of a colonnade in Neville's Court of Trinity College. Lacking anything resembling a stop watch, he adjusted a pendulum to swing in rhythm with successive echoes.⁴² Newton used the pendulum to measure the time it took for a sound to travel from its source down the colonnade to a wall and then back again; he heard the particular sound that he had produced, a moment earlier, reflected back to him. Here we have an example of a sound existing even after the event which produced it ceases (we can suppose) to exist, and which moves independently of whatever produced it. It is also an

⁴¹ This seems to be what David Sanford thinks (1976, pp.192-3). He, like Perkins, takes the relation between sounds and objects with respect to hearing to be analogous to that between colours and objects with respect to sight.

example of a single sound being heard more than once: Newton re-encounters a particular sound, as he must do if he is to time its journey up and down the colonnade. That gives us a reason to think of sounds as a kind object, rather than an event.

Given this, I think that many cases of pure auditory awareness of the source of a sound are best understood as a kind of deferred ostension, as picking out the source of a sound via picking out the sound itself. Thinking of "that car" might be equivalent to thinking something like "the car which is actually producing this sound," of "that dog" as "the dog which is actually making that noise" (c.f. Martin 1997, p.93).

Whether or not one agrees with that suggestion, it remains the case that when one hears an object on the basis of pure auditory experience one does so by hearing the sound that it produces, but one does not hear it as producing that sound. In the kind of cases that I am interested in, we are aware of and can pick out the event that produces the sound independently of hearing that sound, and we can be aware of the event picked out in this way as responsible for the sound we hear.

4. What, if anything, follows from all of this? I began with the idea that sounds can, on reflection, seem not to be part of or fit into the material world of sight and touch. Whilst that may well be true of our pure auditory experience of sounds, sounds heard without using other senses, it is not true generally. We can perceive sounds as being produced by events we see and feel and so to be interacting with the objects of sight and touch. These sounds appear to be as much a part of the material world as the objects we see and feel.

One way to bring out the way in which we ordinarily experience sounds as part of the material world is to consider what our experience of the world would be like if we were deaf. In a passage a little further on from the one I quoted earlier, Proust describes Swann contemplating just this possibility; the deaf man, he supposes, doesn't simply live in silence, the experience of what he sees is different:

⁴² Westfall (1980, pp.455-6).

The highest waterfalls unfold for his eyes alone their sheets of crystal, stiller than the glassy sea, pure as the cascades of Paradise. Since sound was for him, before his deafness, the perceptible form which the cause of movement assumed, objects moved soundlessly now seem to be moved without cause; deprived of the quality of sound, they show a spontaneous activity, seem to be alive. They move, halt, become alight, of their own accord. Of their own accord they vanish into air like the winged monsters of prehistory...the building which the deaf man looks out on from his window...is only so much scenery. If one day it should fall to the ground, it may emit a cloud of dust and leave visible ruins; but, less substantial than even a palace on the stage...it will subside into the magic universe without letting the fall of its heavy blocks of stone tarnish the chastity of the prevailing silence... (p.73).

This seems right: we experience the sounds of the things we see as part of the causal structure of the visible world.

In thinking about how sounds relate to the objects of sight and touch we don't just learn about how we experience sounds as fitting into the world; we can discover something about the nature of the different senses. I have described how we can have experiences of sounds as being produced or caused by things we see; this is a kind of experience which essentially involves more than one sense. When we experience sounds as being produced we have an experience which essentially involves two different senses. In such a case, there are no grounds for thinking that we enjoy two distinct experiences, one associated with vision and the other with audition, so the problem in explaining how we distinguish the senses cannot be that of explaining what the difference between distinct experiences consists in. Nonetheless, on any occasion that we experience a sound as produced by something, there are some things of which we are visually aware, and other things of which we are auditorily aware. That is, we have a perceptual experience of the world, yet in having that experience, we perceive some things visually, and others auditorily. An account of the distinction that we make between the five senses needs to explain in virtue of what that is true. In virtue of what is our perception of

particular things and the properties of things a visual perception, or an auditory perception of them?

Not only does our experience of the production of sounds support my claim about the unity of our experience, it also raises a problem for the Berkeleian and sense-datum theorist's accounts of that distinction. Berkeley thinks that we perceive instances of qualities and that there are different kinds of quality specific to each of the senses. We don't, he says,

immediately perceive by sight any thing beside light, and colours, and figures: or by hearing, any thing but sounds: by the palate, any thing beside tastes: by the smell, beside odours: or by the touch, more than tangible qualities. (1734, 175).

Objects are just collections or congeries of these sensible qualities and there is nothing wrong with the idea that audible ideas should be grouped together with visible ones. The problem raised for Berkeley by the fact that we can experience sounds as produced by objects is to locate some further quality which results from the collection of audible and visible ideas together, but which comes from neither. There's no reason to think that Berkeley couldn't simply add to his list of sensible qualities, so the experience of the production of sounds may not raise any very serious difficulty for this view of experience. But if he were to do so, it would undermine the straightforward explanation that we can give of why we distinguish five different senses.

Since, on a Berkeleian view, experiences are constituted by an awareness of instances of sensory qualities, we can explain why we distinguish five senses by appealing to the fact there are five different kinds of sensory experience, constituted by the awareness of five kinds of sensible qualities. If we add a kind of quality to the list of sensible qualities, then we can no longer explain why we have just five senses by appealing to the fact that there are just five kinds of sensible quality. There will be sensible qualities which cross-cut the distinctions that we make between the senses. If that's right, then we will have to look elsewhere for an explanation of why we distinguish the senses as we do: the Berkeleian view of the nature of our perceptual experience no longer provides such an explanation.

A sense-datum theorist, in contrast, takes our perceptual experience to be constituted by an awareness of mind dependent objects. On this kind of view it is usually supposed that we are not aware of the same particular object with more than one sense, and that the objects of different senses are different in kind. The experiences of different senses are distinct because they are constituted by an awareness of distinct objects, and they are different in kind because constituted by an awareness of different kinds of object. The problem that our experience of the production of sounds raises for this kind of view is that it is an experience of a feature of an object that is both seen and heard. In the face of this the sense datum theorist might simply give up the claim that the experiences of different senses are always constituted by an awareness of distinct particular objects; were they to do so, however, it would no longer be plausible to claim that each of the five senses simply consists in the awareness of a particular kind of object. So again, we can see the kind of experience that I describe as raising a problem for the way that a sense-datum theorist attempts to explain the way that we distinguish different senses. The existence of the kinds of experience that I describe doesn't show that the sense datum view of experience is wrong, just that there is no straightforward way, given that conception of experience, to explain the way we distinguish the five senses.

In this chapter I have concentrated on our experience of sounds, as providing a vivid example of a particular kind of inter-sensory phenomena. Related phenomena occur with the objects of the other senses. When we reflect on our experience of the world we don't discover our experience to be fragmentary, with distinct experiences associated with each sense; just as we experience a unitary world, so our experience of the world is unitary. If I am right then the real problem of the senses is not to explain what makes distinct experiences different, nor to explain how distinct experiences can be related to one another; it is to explain how and why, given the character of our experience, we distinguish five senses in the first place, and to explain what constitutes perceiving some object or property with one sense rather than another. That is the problem that I address in the remainder of this thesis.

4. The Features Account

I have argued that an account of the senses needs to explain what distinguishes the senses. In chapter 2, I argued that this problem is not simply that of explaining how distinct experiences, visual experiences, tactual experiences, and so on, differ from one another. In most cases of multi-sensory perception there are not distinct experiences whose differences we can explain. The problem is, rather, to explain in virtue of what it is true that one's perception of an object is, say, a visual perception of that object; and to explain in virtue of what one's perception of a particular property of an object is an instance of, say, feeling that property. In virtue of what is it true that I can feel the shape of something that I hold in my hand? Remember that in giving an account of the senses I don't just want an explanation of how one can tell with which sense one is perceiving something, but what constitutes perceiving something with a particular sense. It is this question of what constitutes the senses that I will be addressing in this chapter and throughout.

I described, too, how when we reflect on what our experience is like and attempt to describe the difference between, say, seeing and feeling something, what we describe dissolves into a description of differences between what we see and what we feel. This is a consequence of the way we gain introspective knowledge of the character of our experience. When we introspect our experience we attend to just those objects and properties that we attend to in perception; reflection on our experience is, in McCulloch's words, "as world-directed an activity as ordinary seeing itself" (1993, p.53). That makes the fact that we distinguish our perceptions of objects, their properties, and the rest, into instances of seeing, feeling, and so on, seem puzzling. If all we can attend to in introspection is the things of which we are perceptually aware, then what is it about those things which leads us to distinguish our awareness of them into five distinct senses in the way that we do? That is the question that I shall consider in this chapter. The answer, I shall argue, is that there is nothing about the things of which we are aware which will explain the distinction

that we make between the senses and that we must, therefore, look elsewhere for an explanation of that distinction.

If the only differences to which we can introspectively attend are differences in what we are aware of, then we should perhaps look to such differences in giving an account of what distinguishes the senses. We might think, that is, that the senses are to be distinguished by the differing objects or properties that we become aware of by means of them.⁴³ In what follows I begin by describing two problems with this suggestion. The first problem is that it is unable to explain what all the perceptions of a single sense have in common in virtue of which they are perceptions of that sense. The second problem is that it cannot explain, of features or objects that are perceived with more than one sense, what makes a particular perception of such a feature a perception of one sense rather than another. I then go on to discuss a way to elaborate the suggestion to overcome these problems; ultimately, I argue, it fails to do so.

1. Can we explain the distinction that we make between the senses in terms of the *objects* that we are apparently aware of by means of them, and say, for example, that one hears something just in case one perceives a sound? Although perhaps plausible for hearing,⁴⁴ the difficulty with this suggestion is in picking out the objects which are constitutive of each of the other senses. If we list the kinds of objects we can apparently perceive with each sense then, for sight and touch at least, we get a heterogeneous list, as Sorabji, in discussing this kind of view, points out:

there is such a large variety of objects that can be perceived by sight...it would be laborious to define sight by reference to its objects.

⁴³ This is one of the four possible ways in which the senses might be distinguished considered by Grice (1967, p.250).

⁴⁴ Not, one might think, very plausible unless we adopt the Berkleian view that the only direct objects of hearing are sounds. We usually think of physical objects as objects of audition in addition to sounds – ‘I heard Thatcher speak’ one might say. For a more detailed discussion of this kind of view, see Urmson (1968). I want to know whether we can distinguish the senses by appealing to their apparent objects; there are, of course, theories of perception which suppose that there are different kinds of non-apparent objects, such as sense-data, associated with each of the senses.

Moreover it would conceal what unity there is to the concept, and make it a mystery that the single name "sight" should be used to cover such a heterogeneous list (1971, p.60, see also pp.68-9).

It is not just that it would be laborious to define sight in this way, but that appealing to the kinds of objects we apparently perceive will provide no real explanation of what seeing is, and of why we group perceptions together as instances of seeing in the way that we do. We can apparently see a great variety of different kinds of objects: we can see tables and chairs and trees and other material objects; we can see water and glass; we can see the sky and stars and rainbows and the shimmer of hot air. What do all these things have in common in virtue of which we group together all our apparent perceptions of them as instances of seeing? The only thing that they all have in common is precisely that it is possible to see them. If that's true, then we cannot pick out the relevant class of objects, the apparent perception of which is constitutive of seeing, independently of knowing which objects are the possible objects of sight. Something similar is true of the objects of touch. Although we can't tactually perceive all of things that we can see, there are still a large number of different kinds of things that we can perceive tactually. The only thing that these things have in common is that it is possible to perceive them by touch; we have no independent way of specifying the objects perception of which constitute tactual perception.

2. Rather than appeal to the objects of perception, perhaps we should appeal to the features or properties of objects that we apparently perceive in order to distinguish the senses. Such a suggestion has some plausibility since which properties of things we can perceive does depend on the sense with which we perceive them. We might say, for example, that one sees an object just in case one has an apparent experience of it as having some, say, colour or other 'visual' property, that one hears an object just in case one perceives it as having some degree of hardness, etc. If this suggestion is going to work, then we need to be able to say what the properties, our apparent awareness of which is constitutive each sense, have in common other than that we can apparently perceive them by means of that sense.

There is no one kind of property which we can apparently perceive with each sense, but a range of different properties.⁴⁵ We can see things as, for example, having colour, and shape, and a variety of spatial properties – location, orientation, size, and so on – as having textures, and even apparent weight. We can see liquids as having an apparent viscosity; surfaces can look to be hard or soft; tree branches can look flexible; the list is almost endless. There appears to be nothing that all these kinds of properties have in common to which we can appeal in explaining why we group our perceptions of things as having them together as instances of seeing those things. The same is true of touch. The kinds of properties of things that we are apparently aware of by means of touch form a heterogeneous list. In general, the only thing that all the properties that we can apparently perceive by means of a single sense have in common is that they are all perceivable by means of that sense. That means we cannot explain why we classify our apparent perceptions of them as the perceptions of a single sense in a non-circular way. In claiming that the senses are to be distinguished by the properties we are apparently aware of by means of them we will have not given any explanation of why we distinguish the senses in the way we do, we will have simply picked out groups of properties the awareness of which is coextensive with the way that we distinguish the senses.⁴⁶

It might be said, in response to this objection, that there is no reason to require that an account of the senses provide such an explanation. This account of the distinction, it might be claimed, will tell us what is constitutive of each sense, and so what distinguishes each of the senses from the others. Why should we require anything more?

There are two things that we should expect any account of the distinction that we make between the five senses to explain, which the kind of account we are considering – call it the features account – cannot explain. The first is whether other creatures have the same or different senses to those we have, and whether they have a different number of senses to us. We might come across some creature that is able to perceive properties of things that we cannot perceive. It seems

⁴⁵ Taste and smell might be thought exceptions to this claim, but see my comments about taste in chapter 2.

⁴⁶ C.f. Grice (1967, p.255). See also Leon (1988) for objections along the same lines.

reasonable to ask whether it perceives them with a sense which is the same as one of ours, or with a new sense different to any we have. Since, on this account, we have no way of grouping the properties, awareness of which are constitutive of each sense, independently of the way we *actually* distinguish the senses, we must say of such a creature that it perceives with a different sense to ours. But surely the answer is not as straightforward as that. It seems perfectly reasonable to allow that a creature might have the same senses as we do, but that they are more acute or extensive than ours, so that their awareness of the kinds of property in question might constitute an extension of a sense which is of the same kind as one of ours.⁴⁷ Similarly, if we come across a creature which is able to perceive a range of properties quite different to those we are able to perceive, then we would be quite unable, on this account, to decide whether their perception of all these properties constituted the operation of a single sense, or of distinct senses. One might think that is a question to which we ought to be able to provide an answer, and in so far as the features account cannot answer these questions, it is inadequate. We shouldn't, however, rest much weight on these considerations simply because it is not clear how literally we should take our application of our common sense concepts of the five senses to creatures other than ourselves. It may be that the best explanation of our application of these concepts to other creatures is that we do so in a loose, derivative, way which shouldn't be taken literally.⁴⁸

The second thing that an account of our distinction between the senses should explain is why we distinguish just five senses, rather than more or fewer. To do so it must give some account of what is common to all instances of the perceptions of a particular sense in virtue of which they are of that sense. This is

⁴⁷ Perhaps, for a dog, smells have spatial dimensions in a way they don't for us; we wouldn't say, for that reason, that dogs don't have a sense of smell, and instead have some different sense or some sense in addition to the sense of smell.

⁴⁸ Of course, other creatures do have various sensory mechanisms which enable them to perceive their environment, and we can ask whether those mechanisms are relevantly the same as ours; are they, for example, physiologically the same; do they implement the same psychological capacities? It doesn't follow from that that there must be an answer to the question of whether they have the same senses as we do. It would only do so if it were right to identify our five senses with the mechanisms or psychological capacities involved in perception.

something that the features account cannot do. If there is nothing in common to all the properties that we can apparently perceive by means of each sense, then there is nothing about those properties to which we could appeal in order to explain why we distinguish our awareness of them into five kinds, rather than, say, twenty or four. (Equally we will have no way of explaining why the distinction between *five* senses is so widespread.) This point is particularly pressing given that, as I argued, it is puzzling why we distinguish different senses at all. The features account does nothing to solve that puzzle. If we are going to explain the way we group our apparent awareness of different properties together as perceptions of a single sense then we will need to appeal to something other than the properties of which we are aware; we will have to appeal to something like the kinds of sense organ involved, or some feature of the experiences involved.

I have been stressing the fact that the features account is unable to explain what is common to all the perceptions of a single sense in virtue of which they are perceptions of that sense; there is a second kind of problem with the features account. If it is to work, then it needs to be the case that the kinds of properties awareness of which are constitutive of each sense are not the kinds of properties of which we can become aware using any other sense. If we can perceive some property with more than one sense, then it will not be the case that awareness of that property is constitutive of a single sense. Awareness of that property could be the apparent perception of it with any of the senses with which we can perceive that kind of property; such a perception would be of an indeterminate sense. Thus, if the kinds of properties we can perceive cross-cut the way we distinguish the five senses, then we cannot distinguish the senses by appealing to the properties that we apparently perceive by means of them.

There are kinds of properties that we can perceive with more than one sense, the most obvious are spatial properties such as size, and location, motion, and geometrical properties like shape. We can see and feel the shape of something, how big it is, where it is in relation to us, and whether and how it is moving; similarly we can hear where something is, and we can hear it move. And we apply concepts of different senses to other people's awareness of the same kinds of property – we speak, for example, of someone feeling shape of something whose shape they can see. So, if you perceive the shape of something, in virtue of what is

it right to say that you can see its shape? If all we have to appeal to is the fact that you are perceiving its shape, then we will not be able to say whether you are seeing rather than feeling the shape, your perception of the shape would be of an indeterminate sense. Yet it is a determinate matter whether you are seeing or feeling the shape. The features account cannot, therefore, explain the distinction that we make between sight and touch nor, *a fortiori*, between the five senses.

3. It might be said, in response to this, that since we never perceive something as having just a single property – we never, for example, perceive just the shape of something, or just the movement of something – the problem is not what makes the perception of shape a visual perception, but rather what makes the perception of an object as having shape a visual perception. Since when we perceive the shape of an object we also perceive other properties, we can answer that question by appealing to the range of features we perceive something as having. As Dretske puts it:

there is more – much more – involved in seeing an object move than experiencing the object's movement. One also experiences the object's shape, size, color, direction of movement, and a host of other properties (1995, p.95).

The features account can explain what makes the awareness of an object as having shape a visual awareness by appealing to the range of kinds of properties, or the conjunction of determinable properties, of which the subject is aware in being aware of that object. The difference between seeing and feeling the shape of something is that, when you see the shape of something you perceive the shape of the object as one of a range of properties constitutive of visual perception that you perceive that object as having. Had you felt the shape rather than seen it, you would have perceived the shape as one of a *different* range of properties, constitutive of tactual rather than visual awareness. The features account can agree that the same kind of features can be perceived with more than one sense, but claim that, on any occasion, such a feature is perceived as one of a conjunction of features perceived, and that the same conjunction of features could not be perceived with a different

sense. The perception of the properties perceivable with more than one sense will never be of an indeterminate sense, according to this modified features account, because they will always be perceived in conjunction with properties which are determinately constitutive of a single sense.⁴⁹

Appealing to the conjunction of properties of which we are aware will allow us to distinguish different senses, and to distinguish the experiences of particular kinds of property, but only for as long as we restrict ourselves to questions about the operation of each of the senses individually and in isolation from the others. When we see the shape of something our perception of the shape will be a perception of it as having a conjunction of features constitutive of visual experience – that, according to this suggestion, is what makes our awareness of the shape *visual* awareness. But suppose that, as we often do, we simultaneously feel and see an object. Then we would perceive the object's shape in conjunction with kinds of properties constitutive of both visual and tactual awareness: we would perceive the object as having shape, and colour, but also as having weight, texture, and so on. In that case, we can no longer explain what makes the awareness of the object's shape visual by appealing to the fact that we are aware of its shape in conjunction with properties uniquely associated with vision; we are aware of the shape in conjunction with properties associated with both vision and touch. This claim can be made more clearly with the help of an example.

There is a difference between seeing the shape of something – a coin, say – and feeling the shape of that thing; this difference is what constitutes seeing rather than feeling the shape, and what makes true our judgements that we see rather than feel something. An account of the senses must be able to provide an explanation of this difference.

We could deny, as McGinn and Dretske do in the passages that I quoted earlier, that there is any difference between seeing and feeling the coin's shape. They don't deny that what it's like to see rather than feel the shape of the coin is different, but they deny that this difference has anything to do with the experience of the particular property in question – the experience of the shape. According to

⁴⁹ This is in fact the version of the features account that Grice discusses (1967, p.251) see also Roxbee-Cox (1970).

their view, for as long as two experiences represent the very same property – and they don't deny that the experiences of two different senses can do so – then the experiences are, with respect to that property, indistinguishable. But that doesn't mean that they have to deny that there is any difference between the two experiences; there is a difference due to differences in the *other* properties that the two experiences represent. They hold, in other words, a version of the modified features account. When I see the shape of a coin I experience the coin as having shape, size, and colour, my experience represents the coin to have properties it wouldn't represent the coin as having were I feeling its shape. There would be a difference in the properties my experience represents the coin as having, and hence a difference in my experience of the coin's shape. Now, this does provide a description of a difference in what it's like to see rather than feel a coin, and it explains why, if I am feeling a coin what it is like to feel the coin is different to the way it would have been if I were not feeling, but seeing the coin. Suppose, however, that rather than either seeing or feeling the coin I were to simultaneously see and feel it. There are two kinds of case to consider.

In the First Case, suppose that I am holding the coin in the palm of my hand in such a way that I can both feel and see its shape. In the Second Case, suppose that I am holding the coin in the palm of my hand in such a way that, although I can see it's shape, I cannot feel what shape it is.⁵⁰

When I perceive the coin in the First Case is there any difference between the feel of the coin's shape and the look of it's shape? To answer this, consider the difference between the First and Second case. Notice that there is no difference in what I perceive between the First and Second case, there is no difference in what properties my experience represents the coin as having, so there is no difference in

⁵⁰ Grice describes a similar example. "Suppose," he says, "a man be resting a half-crown on the palm of one hand and a penny on the palm of the other: he might (perhaps truthfully) say, 'the half-crown looks to me larger than the penny, though they feel the same size'." If we simply list the features of which the subject is aware we face a problem: "there is nothing in this statement of the facts to tell whether the coins *look* different in size but *feel* the same size, or alternatively *feel* different in size but *look* the same size" (1967, p.253).

the representational content of the experiences; what changes between the First and Second case is not *what* I perceive, just the *way* I perceive it. We might add that, in the Second Case, not only can I tell that I am visually perceiving the coin's shape rather than feeling it, but something about my perception of the shape makes it the case that I am, in fact, visually perceiving the shape rather than feeling it. So we can ask, In virtue of what is my perception of the coin's *shape* visual rather than tactual?

You might think that Dretske, McGinn, and proponents of the features account, can avoid this question simply by insisting that when I both feel and see the coin I am having *two* experiences at one time – a tactual experience and a visual experience – and that these distinct experiences will be different in that each will represent the coin to have a different range of properties as well as representing the coin's shape. Then it would be wrong to say, as I did, that there is no difference in what properties my experience represents the coin as having since in the Second Case, unlike the First, only one of my experiences represents the coin's shape. We could explain the difference between feeling the shape and seeing the shape by appealing to the fact that the shape is represented by an experience which also represents colour, and other visually perceived properties, and that the shape is represented by another experience, apparently distinct from the first, which also represents tactually perceived properties. In this case, what we are aware of in being aware of a difference between seeing and feeling the coin's shape is not a difference in the experience with respect to the property of shape itself, but in the relation between the experience of the property and the experience of other properties.

In order to reply to the objection in this way, the proponent of the features account will need to provide some grounds for claiming that in the First Case I am having two distinct experiences of the coin – a tactual experience representing tactually perceived properties, and a visual one representing visually perceived properties – rather than a single experience which represents both tactually and visually perceived properties. But all they can appeal to in providing such an explanation is our awareness of the coin and its properties, and – as I argued in

chapter 2 – nothing about that awareness provides any basis for the claim that two distinct experiences are involved in our perception of it.⁵¹

The proponent of the features account cannot simply to appeal to the fact that we sometimes talk of seeing and feeling as two different ways of experiencing things, they need to explain what this difference consists in and how it can seem to the subject that she is having two experiences; and, as this example emphasises, that difference cannot be explained merely by appealing to the properties and objects that we apparently perceive.

It seems that in order to explain the way distinguish the senses we need to appeal to more than the kinds of properties that we apparently perceive. Perhaps, then, we should give up the features account and look elsewhere for an explanation of the distinction. Before doing that I want to look at a development of the features account, according to which there are objects or properties associated with each sense which play a special role in the perception of anything else with that sense; such an account claims that the senses can be distinguished by reference to these special objects or properties. One version of it is developed by Roxbee-Cox (1970).

⁵¹ Notice that this objection counts against Dretske's and McGinn's attempt to explain what distinguishes the senses in terms of the representational content of the experiences involved, in terms, that is, of "the properties things are represented as having" (Dretske 1995, p.1). They need to say what makes the representation of the shape of some object, perceived with two or more senses, a visual perception of that shape given that one's experience of the object will represent it as having both tactually and visually perceived properties. To stress again, they cannot appeal to distinct experiences, not least because we have no principled, context independent, way of counting the number of experiences someone is having at any one time. At any waking moment, I am perceiving the world with all five of my senses, I am aware in various ways of my body, and I have various occurrent conscious thoughts. Suppose we were to ask how many experiences I am having. It's not clear what the answer should be. In one sense we can think of someone's experience at any time as constituted by their stream of consciousness and in that sense they only ever have one experience at a time. In another sense we can count the number of experiences which constitute the stream of consciousness as we do, for example, when we ask someone how many pains they can feel. We answer this kind of question by reference to the places the pain is felt, by reference, that is, to the objects of the experience. Neither of these ways of counting experiences provides us with a reason for saying that in the First Case I am having two experiences of the coin.

4. We saw that the problem with the features account is that, when some object is perceived with more than one sense, no explanation can be given of why the perception of, say, the shape of that object is a visual perception. Appealing to the properties perceived in conjunction with shape won't provide such an explanation. But the proponent of the features account needn't simply appeal to the conjunction of properties perceived since, according to Roxbee-Cox, not all the properties we perceive "are of equal standing," instead:

When we perceive something to have a property that [is perceivable with more than one sense] the property will be one the perception of which on any particular occasion requires the...perception of one or other of the special properties I am calling Key Features (p.537).

He claims that when we perceive something to have some property with a particular sense, our perception of it as having that property "involves" or "requires" the perception of it as having some property unique to the sense in question, what he calls the 'Key Feature' of that sense. We can explain what it is to perceive some property with a particular sense by appealing to the Key Feature involved or required for the perception of that property. So, for example, if one's perception of some object involves perceiving that it has some colour or light property – that it has a Key Feature of sight – then one sees that object; and if one's perception that an object has a particular property (other than a Key Feature) involves perceiving that it has a Key Feature of sight then one sees that property. Each of the senses, according to Roxbee-Cox has a unique Key Feature,⁵² and whenever we perceive some object we will perceive it to have some (at least one kind of) Key Feature; if we perceive it as having any other features, our perception of them will involve or require the perception of a Key Feature. We can explain in virtue of what our perception of the shape of an object we see and feel is a visual perception in terms of the Key Feature involved in perceiving the shape: if our perception of the shape

⁵² Roxbee-Cox lists them as, for vision, having some colour; for hearing, having some loudness and timbre; for taste, having some taste; for smell, having some odour; and for touch, having some feel to the touch. He thinks it difficult to characterise the Key Feature of touch and discusses it at some length (1970, Secs. VI and VII).

involves the Key Feature of sight, then we are feeling the shape, and so on. His argument for this is that

[i]f we perceive the shape [of something] by sight, this requires us to...perceive that the thing has some colour property; if we perceive it by touch, this requires us to...perceive that it has some feel to the touch. To suppose that we perceived the shape of a thing, and neither perceived the thing to have some colour property nor perceived it to have some feel to the touch is to suppose something inconsistent with our actual experience of perception...it is therefore natural to say that the perception of shape *involves* or *requires* the...perception of one or other of these properties (p.538).

Is it true that “our actual experience of perception” supports the claim that the perception of one property can require or involve the perception of another. One might think that, at most, it supports the claim that we generally perceive certain kinds of properties of things in conjunction with one another, but if the Key Features account is going to provide an alternative to the features account then it needs to claim more than that an object is perceived to have a Key Feature in addition to other properties that it is perceived to have. We have already seen that that won't distinguish features when an object is perceived with more than one sense, and hence as having more than one Key Feature. It needs to claim more, too, than that the perception of an object as having one kind of property is necessarily or always accompanied by the perception of the object as having a Key Feature since, again, if one perceives the object with two or more senses, one will necessarily perceive it as having two Key Features in addition to whatever other properties one perceives it to have. That will provide no explanation of which properties are perceived with which sense.

Roxbee-Cox recognises this (pp.538-9). He glosses what he means by saying that the perception of one property requires the perception of another, as the perception of one property *through* the perception of another. “I am suggesting,” he says, “that we perceive shapes through...perceiving the presence of some colour property or through perceiving something to have some feel to the touch” (p.539).

But what is it to perceive one property through perceiving another, and how is that supported by our actual experience of perception? Perhaps Roxbee-Cox is claiming that the shape of something would not be perceived if we did not perceive either its colour or its feel to the touch, but not vice versa. If that is right, then it's not just that we perceive a Key Feature in conjunction with other properties, but that our perception of these other properties depends on our perceiving a Key Feature. But what reason is there to think that is true? Someone sceptical of the claim can simply deny it, and Roxbee-Cox offers no further argument in support of the claim.

Nevertheless, the idea that not everything we perceive "is of equal standing" is a common one, and that idea can perhaps developed in such a way as to explain the way that we distinguish the senses. Philosophers of perception sometimes distinguish between immediate and mediate perception, and between corresponding classes of objects which are the immediate or primary objects of perception, and those which are only perceived mediately. David Sanford claims that

[e]verything perceived by a sense having primary objects is perceived by the perception of some primary object (1976, p.194).

He claims, further, that each of the senses has primary objects of different kinds.⁵³ Since everything perceived by a sense is perceived by the perception of some primary object, we can distinguish the senses by reference to them. Something is seen if it is perceived by perceiving a primary object of vision, it is felt if it is perceived by perceiving a primary object of touch, and so on.

What is the contrast between mediate and immediate perception? We might well be sceptical of the distinction. Austin, for example, dismisses the philosophers' use of the notion of immediate perception as not the ordinary use, and says that we are given no definition or explanation of the new use (1962, pp.16-17). We can, however, provide some account of the distinction that is made between immediate

⁵³ The primary objects of hearing are sounds; of touch, tangible expanses; of vision, coloured expanses. Odours are the primary objects of smelling, and tastes are the primary objects of tasting (p.194). This is, of course, reminiscent of Berkeley, except that, for Sanford, the primary objects are physical. Whether our awareness of the primary objects every involves our awareness of anything else non-physical is, he says, another question.

and mediate perception; in what follows I provide such an account, but argue that it will not provide any basis for explaining the distinction that we make between the senses.

Sanford defines the primary objects of a sense as a class of “physical existents” which are such that the perception of any physical existent with the sense in question requires the perception of a member of the class, and that no proper subset of the class satisfies this condition (p.192). He uses our perception of sounds to illustrate what it is for the perception of one object to require the perception of another. Although one can hear something that makes a sound, that thing is not a primary object of hearing because it is possible to hear a sound without hearing the thing that makes it (even if no one ever does hear anything without hearing something that makes the sound). One hears something making a sound, he says, only because one hears the sound that it makes; and anyone who has the capacity to hear, has the capacity to hear sounds (if there are any) which are not produced by anything (p.193).

If this is what it means to say that the perception of one thing requires the perception of another, one might think that the relation is one of causal mediation. Sounds are distinct from the things which cause them; they are caused or produced by those things; and we are auditorily aware of the objects that produce sounds only because we are aware of the sounds they cause. In that sense our perception of sounds is more immediate than our perception of the things that produce them. But this doesn't generalise to the other senses. There is no object which stands to the things that we see and touch as sounds stand to the objects that produce them. Yet Sanford thinks that there are primary objects of vision:

Whenever one sees, one sees either a coloured expanse or light phenomena...Even without agreement on the precise membership of the class of primary objects of sight, we can agree that there is such a class, and that many visible things do not belong to it (p.197).

As he notes, a table does not cause its surface to exist or to have the properties it has, so when we see a table by seeing a part of its surface the table is not seen through the causal mediation of its surface. What then is the relation to which

Sanford is appealing if it's not a causal one? We get a better sense of what he means from the following:

Perception of only primary objects of a sense could be indistinguishable, so far as that sense is concerned, from the perception of things which are not primary objects. The indistinguishability would be due to the sameness of perceptible properties in each case. The primary objects of perception are thus the primary possessors of sensory qualities. Things which are not primary objects have sensible qualities in virtue of their relations to the primary objects which have them (p.194).

Tangible expanses stand to tangible surfaces as coloured expanses stand to coloured surfaces...the perception of just tangible expanses could be indistinguishable, so far as the pressure sense is concerned, from the perception of ordinary solid bodies (p.198).

We can, I suggest, understand Sanford to be making the same distinction as that made by Frank Jackson (1977). Jackson defines the immediate objects of perception in the following way:

...*x* is a *mediate object of (visual) perception* (for *S* at *t*) iff *S* sees *x* at *t*, and there is a *y* such that (*x*≠*y* and) *S* sees *x* in virtue of seeing *y*. An *immediate object of perception* is one that is not mediate; and we can define the relation of *immediately perceiving* thus: *S* immediately perceives *x* at *t* iff *x* is an immediate object of perception for *S* at *t*... (pp.19-20).

The connective 'in virtue of' is not, he says, to be treated as a causal connective, but is rather used to show some analytic definitional relation between the two facts introduced by the terms related. We don't immediately see tables, chairs, and so on, for the following reason. We cannot define the perception of part of an object in terms of the perception of the whole object, because one could see the part without seeing the object had the part been part of another object, and one could see the

object without seeing this part of it if one had seen another part. So seeing an object cannot be necessary nor sufficient for seeing the part, and so cannot be that in virtue of which one sees the part. He concludes that:

We commonly see things in virtue of seeing *other* things: I see the aircraft flying overhead in virtue of seeing its underside (and the aircraft is not identical with its underside); I see the table I am writing on in virtue of seeing its top; I first see England on the cross-channel ferry in virtue of seeing the white cliffs of Dover... (p.19).

One might grant Jackson that we do not see the top of the table in virtue of seeing the table, but deny that it follows from that that we see the table in virtue of seeing its top. Instead, one might claim that we see both the table and the top of the table, and that neither is seen in virtue of seeing the other. How might Jackson defend his claim against this kind of sceptical response? Jackson takes as primitive or fundamental the notions of x seeing y and of x seeing y in virtue of seeing z , and supposes that all the other terms that we use in describing perception and perceptual experience can be defined in terms of the basic distinction between the immediate and mediate objects of perception. His general approach to giving an account of perception prevents him, therefore, from appealing to any other facts about perception to explain why we should accept the contrast.⁵⁴

Nevertheless, there is something that we can say as a way of substantiating both Jackson's and Sanford's claims. In characterising what it means to say that the perception of one thing requires the perception of another, Sanford emphasises that the perception of just the primary objects of a sense would be indistinguishable from the perception of whatever objects we perceive by perceiving the primary objects. This suggests that he has in mind something like the following.

Not all the objects that one perceives on some occasion determine the way that things appear to one on that occasion; there is, one might think, among the objects which one sees at a time, a set of objects which one sees and which look to be a certain way, and which determine how things look to one. For as long as one

⁵⁴ See Jackson (1977, Ch.7) and see Martin (forthcoming, Ch.3, App.2) for a discussion.

continues to see these objects, and they continue to look the same way, then the way things look to one will remain the same, irrespective of what other changes occur to the objects that one sees. Suppose, for example, that you see some array of objects; in perceiving them, they will look to be a certain way. Now, imagine that every part of these objects was destroyed except for their surfaces. Were this to happen, the array of objects would still look to be exactly the same as before: nothing about how things look to you will have changed. But what you perceive will have changed: you will still be perceiving the surfaces of the objects, but will no longer be perceiving the objects themselves; the objects will have been destroyed and will no longer exist.

There is, then, a differential relation between how things look to you, and what you see. Things could not look to you to be how they actually look, had the surfaces of the objects not been as they are and been perceived by you to be so; this is not true of the objects – the non-surfaces – that you perceive. Your experience could appear to be the same as it actually is – things could have looked to be the same way – without you actually perceiving those objects. What is true of vision is true of perception generally. How the things one perceives appear to one to be is determined only by the way parts of the objects that one perceives – their surfaces, or some other part – appear to be. For as long as these parts of the objects remain unchanged, how things appear will remain unchanged, even if one is no longer perceiving the objects of which they are part.

The immediate objects of perception are, on this interpretation, those the perception of which determine how everything perceived is experienced as being. If one also perceives anything else, then that thing is perceived only mediately, by perceiving the immediate objects that one perceives. This is what Sanford means, I suggest, when he says that perception of just the primary objects of a sense would be indistinguishable from the perception of whatever else one perceives by perceiving those primary objects; and it is a way of defending Jackson's claim that we see the table in virtue of seeing its top.

This way of distinguishing immediate and mediate perception relates to Peacocke's distinction between observational and nonobservational concepts.⁵⁵ Peacocke argues that there is a gap between the concepts needed to characterise the way things look to us, and those concepts which can be grounded or applied on the basis of having an experience of things looking a certain way. This is so, he thinks, because something's looking a particular way can be the basis or ground for the application of distinct concepts.

To illustrate this, Peacocke uses the concept of a tomato. We judge that something is a tomato on the basis of its distinctive appearance. But having that appearance is not sufficient for something to be a tomato; to be a tomato something must be of the right biological kind. There are a range of visual appearances shared by real and artificial tomatoes; when we judge that something is a tomato on the basis of how it looks, we do so correctly only if it is a real tomato. An artificial tomato might look exactly like a real tomato, but were we to judge that it was a tomato we would be wrong: the way it appears would be misleading. We can imagine, though, a community who use a concept which applies correctly to some range of tomato-like visual appearances. Their concept would not be the concept of the kind tomato; since it applies to both real and artificial tomatoes it would be the concept, roughly, of a tomato looking thing. Our concept of a tomato contrasts with a concept like square. There is no appearance which square things have, on the basis of which we judge things to be square, which non-square things could have shared. Anything which appears to be square normally is square. We should think of a concept like the concept square as observational because it is the concept of a property whose presence in an object can normally be established by looking from different angles and seeing the object as square. There is, Peacocke suggests, a connection between this distinction between observational and nonobservational concepts and perceptual experience:

perception and observational concepts have to be characterized simultaneously, and one requires for perception matching only in

⁵⁵ See Peacocke (1982, Ch.4) he refines his account in ways that don't affect the point I'm making in (1986, Ch.2); my description of Peacocke and the connection with the mediate/immediate distinction

respect of observational contents. ...when a perceived object is experienced as falling under a nonobservational concept, there must be some level of representational content at which that experience could be perceptual even though the object does not fall under that nonobservational concept. This classifies 'streak produced by the creation of a particle-pair' as nonobservational. It classifies ordinary shape concepts of physical objects as observational. For if components of an experience's representational content containing them are false, there is no more primitive level of representational content which could be true (1982, pp.101-2).

The most primitive level of representational content in Peacocke's account corresponds to Sanford's primary and Jackson's immediate objects of perception. In a situation in which only the surfaces of things exist, our visual experience will be illusory with respect to the tables, tomatoes, and so on, apparently before us. It looks as if there are such objects there, but in fact there are not. One's experience is nonetheless veridical with respect to the surfaces that one sees: one sees the surfaces and they have the properties of shape, colour, and so on, that our observational concepts would ascribe to them. Our observational concepts characterise the immediate objects of perception: in a situation in which there existed solid objects, and not just surfaces, the same observational concepts would be required to veridically characterise the immediate objects of perception.

I have suggested that we can understand Sanford's claim that we see one object by or through seeing another in terms of our judgements about the veridicality of our perceptual experiences relative to the objects of perception in a way that mirrors Peacocke's method for distinguishing observational and nonobservational concepts. Can this distinction between the immediate and mediate objects of perception explain the distinction between the senses and explain in virtue of what we see the shape of something rather than feel it?⁵⁶

owes a great deal to Martin (forthcoming, Ch.3, App.2).

⁵⁶ Note that, as Martin (forthcoming, Ch.3, App.2) points out, in order to draw the contrast between immediate and mediate objects of perception, we must assume that there is an ordering of the

The first question we might ask is whether the immediate objects of perception are the same for both sight and touch. If their primary objects are distinct, then we can explain what it is to see rather than to touch something in terms of the primary objects of seeing and touching. One sees something if one sees it by seeing a coloured expanse;⁵⁷ that is, just in case the primary object of one's experience of it is a coloured expanse. One touches an object just in case the primary object of one's experience of that thing is a tactual expanse. The very same mediate object can be perceived by perceiving distinct primary objects. We can explain, then, what it is to see rather than to touch something; and in virtue of what it is true that you both see and feel the same particular object.

This explanation of what it is to see rather than touch something depends on the truth of the claim that the primary objects of sight and touch are distinct. Were they not distinct, then we could see and touch a particular object by perceiving the same primary object, and we would be unable, in that case, to appeal to a difference in primary objects in order to explain the difference between seeing and touching. So we might ask: are the primary objects of sight and touch always distinct objects? Furthermore, this explanation is an explanation of what it is to see or touch an object, but we also want an explanation of what makes one's perception of some property of an object a perception involving a particular sense. Can we explain in virtue of what one's perception of the shape of something is, say, a visual perception, by appealing to the primary objects of vision?

Are the primary objects of sight and touch always distinct? Sanford claims that they are: the primary objects of touch are tangible expanses, the primary objects

objects of perception which fixes the way that things look to us. We must take it that we perceive the same things in all cases in which things appear to us to be the same way, and that these objects are responsible for things appearing to us to be that way. In the situation in which only the surfaces of things exists then it seems right to say that we perceive the surfaces and that they appear to be a certain way. Is the same true in the situation in which we normally find ourselves, where there are solid objects before us? If not, then the assumption is false. Thomson Clark argues that the conclusion that we perceive only the surfaces of things in the normal case doesn't follow from the fact that that in the first case things appear to be the same way to us, and in that case we perceive only surfaces (1965).

⁵⁷ One sees a red tomato, for example, if one sees an red coloured expanse which is part of a tomato.

of vision are “physically existent coloured expanses”.⁵⁸ But that seems wrong. Given the way I suggested we distinguish between the primary or immediate and the mediate objects of perception, the primary objects of vision are the apparent surfaces of objects. It is the surfaces of things and their properties which determine how things look to us. What about the immediate objects of touch? It would seem that, just as for vision, we can imagine all the parts of the objects we touch except their surfaces having been removed and their still feeling the same. Actually, this may be more difficult to imagine in the case of touch than vision since many of the properties of objects which determine how those objects feel are not properties of their surfaces, hence not properties which could survive unchanged if the objects themselves were destroyed. There are, for example, tactually perceptible properties, such as weight, that are apparent properties of whole objects and not just their surfaces; if all that remained of a piece of fruit, for example, was its surface, then it would no longer feel the same: it would have a different weight and centre of gravity, would squash in one’s hands, and so on. Perhaps, then, it is not possible to imagine a world which appears tactually just like our world, but in which the objects are hollowed out so as to be mere surfaces. However that may be in general, not all our tactual perceptions of objects are experiences as of them having these kinds of property. For at least some tactual experiences, the primary object of the experience – what it is that determines how the object appears tactually – would seem to be the surface of something as, for example, when you run your hand over the top of a table and experience it as a hard flat surface. In this case the table could feel solid even if it were not, for as long as its surface felt solid.

But now, if the primary object of touch in these kinds of cases is the surface of the object perceived, and the primary object of vision is likewise the surface of the object, then we will not be able to say in virtue of what you see rather than touch the object. In both sight and touch you perceive the table by perceiving its surface. So we cannot appeal to a difference in the primary object perceived to explain what it is to see rather than touch the table.

⁵⁸ Or “light phenomena”. He includes light phenomena to cover those cases in which we see something without seeing an apparent coloured expanse: glitters, flashes, and so on.

Although you can certainly perceive the surface of something tactually; we might wonder whether, when you do, there is anything else that you perceive that would determine how things appear to you to be, and yet which is not the apparent surface of an object. It might be suggested that since, when you touch something, you perceive part of your body as it is in contact with the object, it is the apparent properties of part of your body which determine how the object appears to you to be. That doesn't help. Feeling one's body to be a certain way just is for it to appear to one that one is perceiving the surface of something. It's not as though it could appear that one's body was in contact with something and yet it not appear that one was touching something: apparently feeling one's body in contact with an object just is apparently feeling part of an object.⁵⁹ The fact that bodily perception plays a role in touch does not mean, therefore, that we don't immediately perceive the surfaces of objects by touch.

If that is right then the only way to maintain the distinction between sight and touch would be to claim that we never see the surface of things and hence that the immediate objects of sight and touch are never the same. This is just what Sanford does. The primary objects of vision – coloured expanses – are, he argues, not identical with the surfaces of objects:

The description of something as a red expanse does not imply even that it is a surface. A physically existent coloured expanse which was not a surface could be visually indistinguishable from one that was (p.196).

Even if one were not sceptical of the claim that a coloured expanse which was not the surface of an object could be visually indistinguishable from one that was (and Sanford offers no further argument in support of it), its truth does not imply that the immediate objects of vision are never surfaces; to establish that, he would need to argue that there could be a physically existent coloured expanse which is visually indistinguishable from every object or scene we actually see, and not just that there

⁵⁹ C.f. M.G.F. Martin's remarks on the relation between tactual and bodily hallucinations (1992, p.212).

could be some coloured expanse which is indistinguishable in that way, and such a claim would be very implausible. Appealing to the immediate objects of different senses will not, therefore, explain the distinction we make between the senses because both sight and touch can, and on some occasions do, have the same immediate objects.

If the primary objects of sight and touch are the same, then we will not be able to appeal to differences in the primary object in order to explain in virtue of what one's perception of the shape of something is a visual perception of that shape. On this account, one sees the shape of something in virtue of seeing the shape of some part of that thing – its surface, say – and there seems no reason to deny that one could feel the shape of the that object in virtue of feeling the shape of the same part; we cannot, then, appeal to the part of the object one immediately perceives in order to explain why one is seeing the shape.

We can make the same point in terms of Peacocke's distinction between observational and nonobservational concepts. Our concepts of properties like shape are observational, and can be applied on the basis of both seeing and feeling the shape of something. So that we can apply the concept of a shape to something on the basis of some perceptual experience of it does not determine whether or not we are feeling or seeing that shape: we could be doing either. (One might think that there are some concepts for which this is not true: some concepts are only observational when applied on the basis of, say, feeling an object. We can, for example, both feel and see that something is hot, but we can only apply the observational concept to something hot on the basis of feeling something; for, misrepresentation aside, whilst something can look to be hot and not be hot, something cannot feel to be hot and not be.)

5. In conclusion, I have argued that attempts to explain the way that we distinguish the senses by appealing to the objects and features of which we are aware is unsuccessful. A straightforward appeal to the objects and properties of which we are aware won't explain what all the perceptions of a sense have in common; there are, furthermore, features which cross-cut the way that we distinguish the senses. Any attempt to locate features which play a special role in perception also fails: we need to explain in virtue of what one's perception of a particular feature of

something is the perception of a particular sense. Appealing to the other features of which one is aware cannot explain that. We must look elsewhere, then, for an explanation of the way we distinguish between the five senses.

5. The Brute Experiential View

Many philosophers suppose that we can explain the way that we distinguish the five senses by appealing to the character of the experiences involved in perception. A.D Smith, says, for example, that

Necessarily, when we consciously perceive a physical object, that object impinges on us in such a way as to be registered in a sensory state. Such states possess, and are differentiated by, intrinsic experiential features... To perceive a sound as a sound is to be sensorily affected in a peculiar and indefinable way that differs from our sensory state involving, say, the perception of a color. Such sensory characteristics of experience go to define a sense modality (1990, p.239).

Our perceptual experiences, on this view, have intrinsic experiential features which partly determine what it is like to have those experiences. It is in virtue of their having these features that the experience of seeing something is different to the experience of touching something, and so it is in virtue of such features that we distinguish our experiences of things into the experiences of different senses. Similarly, Grice concludes his discussion of the senses by saying that "sight is to be distinguished from other senses by the special character of the experiences involved in seeing," and in general that the senses "are to be distinguished by the special introspectible character of the experiences their exercise involves" (Grice 1967, p.267, and p.259).

I have argued that the only things to which we can introspectively attend are the objects of our experience and that any attempt to describe differences between, say, seeing and touching something, dissolves into a description of differences between what we see and what we touch. I have argued, too, that we cannot explain the way that we distinguish the senses by appealing to the things we apparently perceive. It is for this reason that Grice claims that the special

introspectible character of our experience “resists both inspection and description” (p.259). The way to describe our visual experiences is, he says,

in terms of the way things look to us, and such a description obviously involves the use of property-words. But in addition to the specific differences between visual experiences, signalled by the various property words employed, there is a generic resemblance signalled by the use of the word “look,” which differentiates visual from non-visual sense experience. This difference can be noticed and labelled, but perhaps not further described (p.267).

We can call this account of the way we distinguish the senses the Brute Experiential View, or the Brute View for short. It supposes that there just are brute differences in the character of the experiences of the different senses, and that we can explain the way we distinguish the senses by appealing to these differences. One sees something, on this view, just in case one’s perception of it involves having a visual experience. One sees the shape of something, just in case one’s perception of that shape involves a visual experience.

What reason is there to think that the Brute View is true? I’m going to begin by considering two arguments for the claim that it must be possible to explain the way we distinguish the senses by appealing to the experiences involved. Given that we cannot (as I argued in the previous chapter) explain the way we distinguish the senses by appealing to the objects of our experiences, an explanation of the senses in terms of the experiences involved would have to be in terms of differences in the experiences of the kind appealed to by the Brute View. The conclusion that we must distinguish the senses by appealing to the experiences involved would therefore support the Brute View.

1. One person who apparently holds the Brute View is E.J. Lowe, he says that an

experience qualifies as a visual experience purely by virtue of its intrinsic phenomenal or qualitative character... It is true that such experiences are normally caused by light impinging on the retina of the eye...but it is not

inconceivable that a congenitally blind person (one whose eyes or optic nerves were damaged beyond repair) should be capable of enjoying such experiences and have sight conferred artificially by being fitted with a prosthetic device. The device would not necessarily have to be sensitive to light in the spectral range to which the natural eye is sensitive: indeed, there is no reason in principle why it should not be sensitive to some form of energy other than electromagnetic radiation. So what qualifies an experience as visual has nothing to do with its causal provenance...[q]ualitative character is what counts... (1992, p.80).

We can perhaps all agree that someone could still see even if they had an artificial eye. What makes an artificial eye an artificial *eye*, we might say, is that it serves the same function as a real eye; someone could still see with such an eye precisely because it does serve the same function as a real eye. But it doesn't follow from that that an experience is visual purely in virtue of its intrinsic phenomenal or qualitative character. To draw that conclusion we would have at least to accept Lowe's claim that someone could see with an artificial eye even if it were not sensitive to light. But what reason is there to think that is true? A sceptic might simply deny it.

Perhaps it is a mistake to view the considerations that Lowe offers as an argument for the Brute View, rather than as what amounts to an expression of his commitment to it. His discussion is not after all (and neither is A.D. Smith's), a discussion of how we distinguish the senses. Many of those who accept the Brute View think it obviously true, and so not in need of argument. They think it obviously true that someone could still see in the counterfactual situation that Lowe describes. But we can ask why that should seem so obviously true; and, in order to decide whether or not it is true, we can ask what is the basis of our intuition that someone could still see in that counterfactual situation.

On what basis do we judge that someone in this counterfactual situation could still see? One answer to that question appeals to what we can imagine. Amongst the variety of experiential episodes we enjoy are those of sensory imagining. These are the distinctive episodes of imagining analogous to perception; like perception, we can distinguish different modes of sensory imagining

corresponding to our use of the distinct senses. We can, for example, imagine the view from the hotel looking out over the harbour; imagine hearing a favourite record; imagine the smell of hot coffee; and so on. When we visualise something we imagine how it looks. (There is a similar relation between the other modes of sensory imagining and their corresponding senses.) This suggests that there is some kind of correspondence between the objects of vision and the objects of visualising; there are two different ways in which we might attempt to cash out this correspondence. One is in terms of a correspondence between the objects of visual experience and visualising: the way I visualise things is the way that they would look if veridically perceived. The other way is by appealing to visual experience: so, we might say that to visualise some object is to imagine seeing – having a visual experience of – that object. The first claim is weaker than the second in that it sees the correspondence between visualising and vision as following from the fact that we can imagine the same things as we can perceive, but not that we imagine things by imagining perceiving them. In what follows, I am going to suppose that we can, at least sometimes, imagine seeing something.

We can take the fact that when we imagine, say, seeing something, we imagine having a possible visual experience, to be evidence for a certain possibility: that we can imagine such an experience is evidence that we could actually have an experience of the kind we imagine having. If it is true that when we imagine seeing something we imagine a possible visual experience, then what we imagine when we imagine seeing something might be thought to provide evidence about the circumstances in which it is in fact possible to see something, hence to provide evidence for the kinds of counterfactual claim that Lowe makes. If that line of reasoning is correct, then we can ask what we must imagine in order to imagine seeing something, and take an answer to that question to be evidence for what must be the case in order to actually see something.

When we imagine seeing something we imagine how that thing looks, we imagine something as having certain kind of appearance. That what we imagine has that appearance would seem to be sufficient for our episode of imagination to be an imagined seeing of something. In order to imagine seeing something we don't have to imagine anything other than how that thing appears; in particular, we don't have to imagine anything about the causal provenance of the experience we imagine

having. What we imagine when we imagine seeing something is neutral on the question of how that experience was produced, what kinds of mechanisms – physiological or psychological – were involved, whether the radiation which mediated our awareness was electromagnetic or of some other kind, and so on.

What can this kind of imaginative project tell us about the nature of our actual experience? It can tell us about what experiences it is possible to have. If what we imagine when we imagine seeing is a possible experience, then it is possible to actually have an experience of kind we imagine having.

All I needed to imagine in order to imagine seeing something is things having a certain kind of appearance. That I imagine things looking a certain way is sufficient for me to imagine seeing; I don't need to imagine anything else – I don't need to imagine the experience as having any particular kind of causal provenance – in order to imagine seeing. We might conclude that, since in order to imagine that experience I didn't need to imagine anything other than a certain kind of appearance, it is possible for me to actually have an experience of the kind that I imagine just by having an experience of things appearing a certain way and irrespective of the causal provenance of that experience. Such an experience would be visual experience.

If that is right, then it is possible for me to have an experience of the same kind as I imagine having – a visual experience – just by having an experience of things appearing a certain way. The causal provenance of the experience doesn't matter; its causal provenance is not necessary for an experience being the kind of experience it is. The same is true of the other senses. So we can conclude that all that is necessary in order to have an experience of a particular sense is that we have an experience with a certain qualitative character, and all that needs to be the case in order to perceive something with a particular sense is that our perception of it involves an experience with that kind of qualitative character. Something like this kind of reasoning, or at least reflection on the kinds of experience we can imagine having, is what grounds our intuitions about counterfactual situations of the sort described by Lowe; and it is such reasoning which is the basis for the conclusion that all that matters is the qualitative character of an experience: that the causal provenance of an experience doesn't matter – that the eyes, kind of mediation, and the rest, play no essential role in seeing.

Should we accept the conclusion of this line of reasoning? I suggest not, for two kinds of reason. The first has to do with what it is to imagine a visual appearance. When we imagine the visual appearance of something we imagine it to have certain kinds of properties – those visually perceptible properties which determine the way something looks. This is essentially so: you cannot imagine the visual appearance of something except by imagining it to have a certain range of visually perceptible properties which determine it as having the appearance it has. That limits the kinds of things that you can imagine seeing. You cannot, for example, imagine seeing sounds, since sounds lack visually perceptible properties. It might be objected that we can certainly visualise sounds. That is not a counterexample to my claim because to visualise a sound is just to imagine it as having a certain kind of visual appearance, that is, as having visually perceptible properties. We might question, too, whether it is possible to imagine seeing something without imagining seeing light, or at least, illumination and illuminated surfaces. Light contributes to the way things visually appear. Think, for example, of all the light phenomena that we perceive and which contribute to the visual appearance of things: shadows, colours, reflections, and so on. The visual world is suffused with light.⁶⁰ If our visual experience had been mediated by something other than light, then things would have appeared differently. Perhaps we can imagine seeing mediated by some kind of radiation other than light, but what we imagine, when we imagine seeing that, is something which appears in the way that light appears. So the argument only gives us a reason to think light is not necessary for seeing if there could be something that appears like light, but isn't light.⁶¹

The argument doesn't, therefore, give us any reason to think that it is *merely* the qualitative character of an experience that matters to its being the kind of experience it is because it gives us no reason to think that we can identify visual experiences independently of identifying them as experiences of certain kinds of visually perceptible properties. That we cannot identify visual experiences independently of identifying them as experiences of certain kinds of properties leads

⁶⁰ For a discussion of the role of light in seeing, see O'Shaughnessy (1985).

⁶¹ That suggests that Kripke is wrong when he says that we can imagine that sound waves gave some creature "visual impressions just as we have" (1980, p.130).

Roxbee-Cox to reject the idea that the distinction we make between the senses can be explained in terms of the experiences involved. Such an explanation of the distinction would, he says, allow for the following possibility:

The visual experience, that is, one that normally accompanies our perception of, for example, colours, might on some occasion accompany the perception of sounds, and we should have a situation of sounds being seen. If such an eventuality is to be describable, then visual experience must itself be identifiable other than as whatever experience one has when one sees (1970, p.533).

That, I have suggested, is not a genuine possibility. The Brute View, as characterised by Grice, isn't committed to this possibility. It is consistent to hold that experiences can differ in virtue of having features which we cannot identify independently of the objects and properties we are aware of in having those experiences. We might still wonder, if the properties of the things we imagine play an important role in explaining what makes an appearance visual, then what reason is there to think that we need to appeal to anything else in order to explain what makes an imagined experience a visual experience? I consider that question in the next section.

The second reason for rejecting the conclusion of the argument from imagination comes from questioning the kind of evidence that an episode of imagination can provide about what may be actual. The point here is a simple one. When you imagine experiencing something you imagine having a possible experience, an experience, that is, which it would be possible for you to have; so we can take your imagined experience as evidence of what is possible in the same way that we take actual experience to be evidence for what is actual. If this is how an imagined experience can provide evidence for what is possible, then it is hard to see how our imagining a visual experience can tell us whether or not the causal provenance of the experience is necessary for it to be the kind of experience it is: nothing will be revealed by an imagined visual experience about its nature that is not revealed by an actual visual experience. If an actual visual experience doesn't reveal whether or not its causal provenance is necessary, then how could an imagined

experience help? What we imagine is an experience just like the actual experience, and if the actual experience doesn't reveal the nature of the experience, then an imagined experience just like the actual one won't either.

Often acts of imagination have content that goes beyond what is present in their imagistic content. You might imagine a ripe tomato by imagining the red colour of its skin; but you could visualise a ripe tomato without visualising its colour. In doing so you would still imagine something red. The difference between the two episodes of imagining is in the presence or absence of colour in the sensory content of what is imagined, but the object of imagination – a ripe tomato – is the same in both cases. The same sensory content can be put to different imaginary purposes. This extra non-sensory content Peacocke calls the 'S-imagined' conditions of an act of imagination, where S-imagining "shares with supposition the property that what is S-imagined is not determined by the subject's images, his imagined experiences" (1985b, p.25). When we imagine changes in the causal provenance of an experience – when we imagine the presence of certain mechanism, say – we don't imagine changes in the sensory content of what we imagine. The difference between imagining a visual experience as being caused in one way rather than another, or not being caused at all, is not a difference in the image itself, not a difference in the way we imagine the experience to appear; it is, rather, a difference in the S-imagined conditions that accompany the image. So to the extent we think such an imagined experience is a reliable guide to what is possible it must be because we regard that what we can S-imagine to be a reliable guide. But it was to answer the question of whether our supposition that someone could see irrespective of the causal provenance of their experience that we appealed to what we can imagine in the first place. What we can imagine, then, is not going to provide any reason for or against that question (c.f. Peacocke 1985b, pp.33-34). That we can imagine having a visual experience without imagining it as produced or caused in some particular way cannot show that it is possible to have a visual experience which is not produced or caused in some particular way. What makes this a case of imagining having a visual experience which is not caused in any particular way is the accompanying conditions – the S-imagining – and nothing about the imagined experience itself.

The argument that I described, from what it is possible to imagine, doesn't give us any reason to think that the causal provenance of an experience doesn't matter to its being the kind of experience that it is, and so it give us no reason to think that qualitative character of an experience is all that matters to its being, say, a visual experience. It gives us no reason, therefore, to think that the senses can be distinguished by appeal to experience alone, and so gives no support to the Brute View. In the next section I consider another way in which the fact that we have different modes of sensory imagination might be thought to support the Brute View.

2. In the previous section I asked whether the fact that we can imagine seeing something independently of imaging its causal provenance gave us a reason to think that the Brute View is true. In this section I am going to consider a slightly different question: Does the fact that we enjoy different modes of sensory imagination give us any reason to think that an explanation of how we distinguish the senses must appeal solely to the kinds of experience involved. We might think that it does if the only possible explanation of these different modes of imagining was one that appealed to the character of the experiences involved in perceiving.

I described how episodes of imagining things to be a certain way can have both sensory and non-sensory aspects. In what follows I shall suppose that the difference between the different modes of sensory imagining is a difference in their sensory content. So that, for some cases at least, we can say what makes a visualising a visualising rather than some other kind of sensory imagining simply by appealing to some aspect of the sensory content of that episode of imagining. Robert Hopkins has argued for an account of visualising – visual imagining – and its relation to vision which, if correct, suggests that we can distinguish the senses merely by appealing to the kinds of experience involved (1998, Ch.7). I will begin by examining his account.

There are two different ways in which we might attempt to cash out the correspondence between the objects of vision and the objects of visualising: in terms of a correspondence between the objects of visual experience and visualising, or by appealing to visual experience. Hopkins' account is of the first kind, he explains the correspondence in the following way: "[e]very visualising has some

content which matches *part* of the content of a possible visual experience" (p.169). The idea here is simply that part of what is visualised is necessarily some of what, in a possible visual experience, could be seen. This, Hopkins says, "entails that visualising always ascribes a visual appearance to its objects, that it represents them as looking a certain way" (*ibid.*). He suggests, furthermore, that its doing so is what distinguishes visualising from the other modes of sensory imagination: to visualise something is to imagine it looking a certain way, to imagine touching something is to imagine it feeling a certain way, and so on.

If this suggestion is to be at all illuminating we need to say what it is to imagine something as looking a certain way and how that differs from imagining it as feeling a certain way. In claiming that visualising is that class every member of which shares some content with visual experience, we are supposing that we have already distinguished episodes of imagining into classes – visual, tactual, and so on – given that we have done so, we can then identify the different classes in the way suggested. But we need to say something about how the various episodes of experiential imagination are to be divided into classes corresponding to each sense. What is it that all members of the class of, for example, visualisings have in common? Hopkins calls this "the classification problem" (p.170). The obvious solution to this problem, he suggests, is to identify the members of each class in the same way that we do the class itself. But there is a problem with doing this. There could be an episode of visualising with a content which matches that of some non-visual experience as well as that of some visual experience. If the account is to explain what makes this episode of imagining a case of visualising, it needs to exclude this possibility.

He thinks that we can explain what visualisings all have in common "by specifying certain aspects of their content which [visualising] and seeing necessarily share" (p.171). But merely appealing to the kinds of properties represented by the experiences of different senses and by different modes of imagination won't work. There don't appear to be any properties which are such that all and only visual experiences represent them and which are necessarily represented by visualisings. So, rather than appeal to the sorts of properties represented, Hopkins suggests we should look instead to the way that vision and visualising represent those properties;

in particular we should look to the fact that both represent properties in a way that is perspectival.

To claim that an experience represents things in a way that is perspectival is to claim that whatever is experienced is experienced as presented to a point, such that one experiences things as oriented relative to that point. Hopkins claims that vision and visualising are always perspectival in this way. We can always answer questions like 'Which way round is it?' about the things that we visualise, just as we can for the things that we see. We cannot appeal to this fact alone, however, in order to distinguish visualising from the other modes of sensory imagination because both touch and vision are perspectival. They both "represent space, the objects which occupy it, and the spatial relations between those objects; but, more than this, they represent that space from a point within it" (p.172). So the content of visualisings share with both visual and tactual experience the property of being perspectival. There is, nonetheless, a difference in the way that vision and touch are perspectival. They differ, Hopkins claims, in how they represent objects in space and the relation of those objects to the point from which they are perceived. The difference arises as a consequence of the fact that we must bring some part of our body into contact with something in order to perceive it by touch. In touch we are aware of the spatial location of parts of our body, and of sensations of touch where parts of our body come into contact with the objects we perceive. One perceives an object tactually in virtue of being aware of a single place as both the location of the sensation of bodily contact, and as part of the object that is in contact with one. In touch, then, some part of the perceived object is always presented as in the same place as some part of one's body. This contrasts with vision. In vision, things are always presented as spatially separate from the point to which they are presented, the point of view. Whatever is seen is seen as occupying a place separate from but related to one's point of view; touch never involves this sort of perspective, so we can explain what is distinctive of visualising by appealing to this feature of the way things are represented.

If Hopkins is right then we can explain what is characteristic of visualising in terms of a perspectival feature of its content it shares with the content of visual experiences. But in giving an account of visualising he has also given us an account

of what makes an experience a visual experience as distinct from a tactual experience.⁶²

Whatever the merits of his explanation of sensory imagining, Hopkins' account of what makes an experience visual or tactual can't explain the way that we distinguish the senses. It can't do so because in certain cases it cannot explain what makes our awareness of some property of an object a visual rather than a tactual awareness of that property.

A tactual experience might be thought to be capable of representing some part of an object, presented as occupying the same place as part of one's body, as a point on the rim of a glass, or as part of the surface of a cubic object, or as a place on a flat surface. In claiming that tactual experience always represents some part of the perceived object as in the same place as part of the subject's body, Hopkins is not committed to denying this. He doesn't claim – and surely he is right not to – that the content of a tactual experience is exhausted by the representation of such points of contact. We are tactually aware of objects and their features in virtue of touching only some part of them. His claim, then, is not that, for any part of an object, if that part is to be represented by a tactual experience, then it must be represented as in the same place as a part of the subject's body in contact with it. But rather, that if a tactual experience is to represent some object, then it must represent some part of that object as in the same place as some point of contact. Given this, he claims that even the most generous account of the content of tactual experiences could be squared with his account of the perspectival character of such experiences (p.178).

Consider how, then, on Hopkins' account, we might explain what it is to feel the shape of some object that we perceive. It is to have an experience which

⁶² What about the other senses? Can the same feature be used to distinguish the experiences associated with the other sensory modalities? Hopkins suggests that hearing is perspectival in the same way that vision is, with the following difference. The objects of hearing, unlike those of vision, may not be presented as spatially separate from the point to which they are presented. He is not sure how to describe the categorical basis of this modal difference, but he nonetheless remains convinced that we can distinguish the experiences involved in this way. We cannot extend the account to smell and taste (pp.181 ff.), but it's more plausible to appeal, in these cases, to the features represented by the experiences.

represents the shape of an object some part of which is represented as in contact with some part of one's body. This contrasts with vision. To see the shape of an object is to have an experience which represents the shape of an object which is represented as spatially separate from one's point of view.

Suppose, though, that you see, but don't feel, the shape of an object that you both see and feel. How might we explain in virtue of what your awareness of the object's shape is a visual awareness of that shape. Appealing to the content of your experience doesn't seem sufficient: you have an experience which represents the shape of an object which is represented as spatially separate from your point of view and part of which is represented as in contact with a part of your body. From this description of the content we can conclude that you are both seeing and touching the object, but there is nothing in the description to tell us whether you see the object's shape or feel it.

Hopkins describes his account as an account of the different ways that properties are represented in vision and touch, but really the only difference in the way that properties are represented is a difference in the way that the objects which have those properties are represented. This is clear in the case of touch. His claim is that some part of every object one touches is represented as in the same place as a part of one's body. That makes the representation of that object a tactual representation. We cannot explain what makes a representation of a property tactual in this way: one doesn't touch parts of properties, but parts of objects that have properties. So one's experience of a property must be tactual in virtue of its being represented as a property of an object one touches. So the relation between one's tactual awareness of the shape of an object and one's awareness of the points of contact with that object is that one is aware of the shape of an object that one touches. That is not sufficient to explain why the awareness is tactual when it is also an object that one sees. That it cannot do so doesn't necessarily mean that it is inadequate as an account of what distinguishes the different modes of sensory imagination. Whether or not it is so is a further question. Rather than pursue that question directly, I want to ask a slightly different question.

Does the fact that we experience these different modes of sensory imagination imply that there must be different kinds of experience associated with

each of the senses, and hence that an account of what distinguishes the senses need appeal only to the experiences involved?

How might an argument for that conclusion go? The idea, roughly, is this. An explanation of how we distinguish episodes of sensory imagination into different modes can appeal only to features of their sensory content, features shared with perceptual experiences; if these features are, as they must be, sufficient to explain how we distinguish modes of imagination, then they must be sufficient to explain how we distinguish perceptual experiences into kinds corresponding to the senses. If features of sensory content are sufficient to explain how we distinguish perceptual experiences into such kinds, then nothing else is necessary; in particular, no external factors are necessary.

It is undeniable that there are different modes of sensory imagination corresponding to our different modes of sensory perception – visualising, imaginary touching, imaginary hearing, and so on. I said, furthermore, that I am going to suppose that these modes of imagining are (or, at least, can be) different in virtue of differences in their sensory content alone. That means that we can say what makes an episode of imagining a visualising rather than, say, an episode of tactual imagining, simply by appealing to its sensory content. It is not necessary to appeal to the wider cognitive project of which that episode forms a part.

Given that supposition, there must be something about the sensory content of an episode of sensory imagining in virtue of which it is an episode of, say, visualising rather than some other mode. There must be some feature of all members of the class of visualisings which determine them as members of that class; if the only difference between an episode of visualising and an episode of tactual imagining is a difference in its sensory content, then what unites all members of visualising must be some feature of that sensory content, some experiential feature. In general, there must be some feature of the sensory content of episodes of sensory imagination sufficient to determine them as episodes of a particular mode of sensory imagining.

Not only are the differences between modes of imagination sensory differences, but we think there is a correspondence between these different modes and our modes of sensory awareness; this correspondence is a correspondence which holds in virtue of some shared feature of the sensory content of the states. It

seems plausible that whatever feature of the sensory content of an episode of imagining determines it as of a particular mode would be sufficient to determine a sensory experience as an experience of that mode. If there is a feature of the sensory content of an experience sufficient to determine it as an experience of a particular sense, then nothing else, in particular no external factor, is necessary to so determine it. If the difference between episodes of imagining is a sensory difference then the difference between modes of sensory experience must be a sensory difference too.

If that's right, we should conclude that there is some feature, shared by both sensory experiences and episodes of sensory imagination which determine them as episodes of a particular mode – vision and visualising, and so on. So, even if Hopkins is wrong in his characterisation of the feature in question, there must be some feature of experience that will distinguish the different modes of sensory imagining and sensory experiences. We should conclude, too, that we can explain the way we distinguish different senses by appealing to the experiences involved: there is some feature of the sensory content of all visual experiences in virtue of which they are visual experiences and which is sufficient to explain what makes a perception of something an instance of seeing that thing.

It looks like the fact that we have different modes of imagination which must be explained in terms of some difference in their sensory content implies that external factors play no role in an explanation of what distinguishes the senses. That's a conclusion that I want to resist. I think that external factors do have a role to play in explaining that distinction. For such a claim to be plausible I need to show how it is consistent with an explanation of the different modes of sensory imagining; in particular I need to show that the existence of such distinct modes of imagining doesn't imply that there are distinct kinds of experience associated with each of our modes of sensory awareness sufficient to distinguish them. I don't deny either that there are different modes of sensory imagining, or that the difference between different modes is solely a difference in their sensory content; but I want to suggest that that is consistent with denying that we distinguish the senses by reference to the kinds of experiences involved in perception.

An account of how we distinguish the senses needs to say what all visual experiences have in common in virtue of which they are visual experiences. If it

cannot do that then it will not be able to explain why all visual experiences are the experiences of a single sense modality. The argument that I have been considering claims that there must be some feature of the sensory content of these experiences in virtue of which we can group them together as experiences of the same sense, that experiences associated with each sense form an experiential kind.

The argument depends on the assumption that whatever feature of the sensory content of a visualising which is necessary to explain why it is a visualising would be sufficient to explain what is it that makes a visual experience visual. And in general, whatever feature of the sensory content of a sensory imagining is necessary to determine it as the mode of imagining it is, would be sufficient to determine the corresponding mode of sensory experience.

Could we deny that assumption and, if so, why? Every episode of visualising must have some feature which relates it to vision and distinguishes it from touch and the other senses. For both sensory imagination and perceptual experience we need to explain why we make distinctions between different senses; we need to explain why we distinguish episodes of imagination and perceptual experiences as we do. But we have different materials available for an explanation in each case. Given that we distinguish the senses as we do, the problem with explaining imagination would seem to be that of explaining why we distinguish episodes of imagination into modes *corresponding* to the senses; but we needn't suppose that we could distinguish episodes of imagination as we do independently of an explanation of the way that we distinguish the senses. So we needn't suppose that all visualisings, say, have anything in common except their relation to visual experience. So although every episode of visualising must have some feature which relates it to vision and distinguishes it from touch and the other senses, it's not the case that there must be some one feature that every visualising has which relates it to visual experience. That means that visualisings don't necessarily form a kind; hence that they don't necessarily form a kind in virtue of some feature of their sensory content.

If visualisings don't necessarily form an experiential kind, then the fact that we distinguish different modes of sensory imagination gives us no reason to suppose that there is some sensory feature common to all visual experiences in virtue of which they form an experiential kind. If visual experiences do not form an

experiential kind then we will not be able to explain why all visual experiences are the experiences of a single sense by appealing to the kind of experience involved.

My suggestion is, then, that what unites members of the class of visualisings need not be intrinsic to the sensory content of visualisings. But if there is nothing common to the sensory content of all visualisings then how do we explain visualising: how do we explain what distinguishes visualising from the other modes of sensory imagination?

We can say what is constitutive of the different modes of sensory imagination just in case we can say why we relate them to the corresponding modes of perceptual experience in the way that we do. We can do that without supposing that all the episodes of a particular mode of imagination have something intrinsically in common. We might say, for example, that an episode of imagining is a visualising if it shares some feature or group of features with visual experiences which it doesn't share with tactual experiences, or because it lacks some feature which tactual experiences necessarily have. Since all we have to explain is a correspondence between imagination and perceptual experience, we needn't suppose that whatever features explain that correspondence would be sufficient to explain the way that we distinguish different senses.

I am denying that there is anything common to all visualisings, and have suggested that we can, consistently with that, still give some account of what distinguishes visualising from tactual imagination. Is there any reason to think that visualisings must form a kind independent of their relation to perceptual experiences? Perhaps there are. You might think, for example, that when you visualise something there's just a single kind of thing that you do; or you might think that there is some underlying mechanism responsible for all visualisings and in virtue of which they form a kind. Even if true that wouldn't show that visualisings form kinds in virtue of their sensory content, and so would give us no reason to think that perceptual experiences form experiential kinds corresponding to the different senses.

In conclusion, then, the fact that we distinguish modes of imagination corresponding to the five senses doesn't imply that we can explain the way we distinguish the senses by appealing to the experiences involved, and so doesn't provide us with a reason for thinking that the Brute View must be true.

3. I have considered two arguments which aim to show that we must be able explain the differences between the senses by appealing differences in the experiences involved. Neither of these arguments is compelling, but that does not, of course, mean that the Brute View is false. Is there any reason for rejecting it?

There has been a tendency of late for some writers to deny that experiences can differ from one another in virtue of having features of the kind the Brute View appeals to. Such writers claim that the differences in what it is like to have an experience can be completely characterised in terms of differences in their representational content, in terms, that is, of what they are experiences of. Someone who makes this claim accepts what Peacocke calls the 'Adequacy Thesis' – the thesis "that a complete intrinsic characterisation of an experience can be given in by embedding within an operator like 'it visually appears to the subject that...'⁶³ some complex condition concerning physical objects" (1982, p.8). Peacocke argues that we should reject the Adequacy Thesis on the grounds that there can exist a "a pair of experiences...[which] have the same representational content, but differ in some other intrinsic respect" (p.13). Those who accept the thesis deny that experiences ever differ in the way that Peacocke describes. Dretske, for example, thinks that "how things seem to us at the sensory level, is constituted by properties things are represented as having" (1995, p.1).⁶³ The difficulty in deciding whether or not the Adequacy Thesis is true (and part of the reason, perhaps, for the continuing disagreement about the existence of such features of experience) is that, in order to decide its truth, we need to be able to decide whether two experiences, which differ in what it is like to have them, also differ in their representational content; and, at the moment, there is no agreed way of determining when two experiences have the same representational content.⁶⁴

If one accepts the Adequacy Thesis then one will think that experiences do not have the features of the kind appealed to by the Brute View in explaining the way that we distinguish the five senses; hence one will reject the Brute View. I don't

⁶³ Many others agree with Dretske. See, for example, Harman (1990), Lycan (1996), and Tye (1992). It is not clear whether Peacocke would still claim that all the examples he described in (1982, Ch.1) tell against the Adequacy Thesis; see Peacocke (1992, Ch.3).

want to take sides in the debate over the existence of such features of experience appealed to by the Brute View. Instead I want to ask whether, even if we accept that experiences can have such features, what reason is there to think that they explain the distinction we make between the senses?

Grice thinks that we should accept the Brute View because it provides the best, indeed the only, explanation of the way we distinguish the senses. He asks how we could decide a "claim to the effect that certain creatures possess a faculty which should be counted as a sense, different from any of those with which we are familiar?" (1967 p.248). To answer this question we will need to decide how we in fact distinguish the familiar senses, and then determine whether these creatures have the same or different senses to ours. He considers various explanations of how we in fact distinguish the familiar senses, none of which is adequate. The Brute View is, it seems, the only plausible explanation, indeed the only possible way to answer the question.

But we might wonder how good an explanation the Brute View can really provide. The first thing to notice is that, for the Brute View to be true – for it to provide an explanation of the distinction – it is not sufficient that our perceptual experiences have the kinds of features to which it appeals. We could accept that experiences have such features and still deny or doubt the truth of the Brute View; doubt, that is, that the distinctions we make between different senses can be explained by appealing to such features. Why would we do that?

For the Brute View to provide an explanation of the distinction there must be five different kinds of features of experience, corresponding to the different senses. It is not enough that the way something looks should differ from the way it feels in virtue of our perceptual experiences having these characteristic features; it also has to be the case that all the perceptions that we classify as visual involve experiences with the same kind of feature: that is, we need to explain what all the perceptions of a particular sense have in common in virtue of which they are perceptions of that sense. The Brute View claims that all visual experiences are alike in sharing the same kind of feature. But even if we grant that experiences can differ in virtue of having such features it doesn't follow all the experiences of a

⁶⁴ For a discussion of this see Martin (1998b, p.167).

particular sense share the same kind of feature. If, for example, there are more than five kinds of feature that experiences can have, then appealing to those features won't explain what all the perceptions of a particular sense have in common, and won't explain why we distinguish five senses. It doesn't follow, then, from the fact that experiences differ in this way, that such differences are what constitute the differences between the senses.

On what grounds might someone claim that all the experiences of a particular sense, all seeings for example, have some feature of experience in common? It doesn't follow, of course, from the fact that what it is like to see something is different from what it is like to feel something. In many cases it may be that we can explain why what it is like to see something is different from what it is like to feel something simply by appealing to differences in what is seen and felt. But even when that is not true, and we explain the difference by appealing to the features of the experiences involved, it doesn't follow that all instances of seeing something share the same kind of feature, and it doesn't follow either that such differences in what it is like to have the experiences constitute the difference between the senses. It might be suggested that the fact that we group the experiences of particular senses them together and call them experiences of seeing and experiences of touching gives us a reason to suppose that they all have something in common. That is, the fact that we distinguish them into experiences associated with each of the senses gives us a reason to suppose that all the experiences of a particular sense have something in common. But that just begs the question: unless we can provide no alternative explanation of the distinction, the fact that we group our experiences as we do doesn't give us a reason for thinking that there are five kinds of features of experience. Nor will appealing the fact that there is something similar in what it is like to have experiences of a particular sense. Someone sceptical of the Brute View can claim that, to the extent that all instances of seeing something seem similar to one another, we can explain that by appealing to the fact that many instances of visual perception will be perceptions of similar kinds of objects and properties, and that that is enough to explain why the

experiences seem similar.⁶⁵ The same thing could be claimed about the sense of touch, and the other senses. We can explain the similarity of the experiences of different senses without appealing to features of the experiences.

I have not argued that the Brute View is false. I have suggested that our reason for accepting it is that it is the best explanation of the distinction we make between the senses, and I have suggested that it does not provide a very satisfactory explanation. It doesn't do so because we have no independent reason for thinking that the conditions it must meet in order to provide such an explanation – that there exist five kinds of feature that experiences can have – are met. Given that, if we can provide an alternative explanation of the way we distinguish the senses, we will have reason to reject the Brute View. I provide such an explanation in my final chapter; I will return there to the question of whether it can provide a better explanation of the distinction than the Brute View.

⁶⁵ This is, of course, exactly what Dretkse, and the others who deny the existence of such features, claim. What it is like to see rather than touch something does not, therefore, provide unequivocal support for the Brute View.

6. A Simple Theory of Perception

So far I have been looking at attempts to explain the distinction that we make between the five senses which appeal to something about our perceptual experience: either to the things that we are apparently aware of, or to features of the experiences in virtue of which we are apparently aware of those things. It might be thought, however, that we can explain how and why we distinguish the senses by appealing to some factors 'external' to our perceptual experience, factors such as the involvement of the sense organs, for example. In this chapter and the next I consider that suggestion. We can distinguish between those common-sense or everyday external factors which can be grasped by everyone, and those which can only be uncovered by scientific investigation. I begin, in this chapter, by considering the former. I discuss the latter in the next chapter.

If the distinction that we make between the senses is made as a consequence of the role that certain external factors have in our perception, then, in grasping the distinction, we must understand how those factors contribute to distinguishing the senses. We can think of that understanding as involving a theoretical understanding of the ways that the production of different perceptions involves various different external factors. It is sometimes claimed that, in understanding our perceptual experience as the experience of an independent world, we must grasp a simple theory of perception. In this chapter I consider whether our grasp of such a simple theory of perception could explain our grasp of the distinction we make between five senses. I begin by elucidating the idea of a simple theory of perception, and suggest what conditions must be satisfied if our understanding of the distinctions is to consist in grasp of such a theory. I then argue that these conditions are not satisfied, and conclude that everyday external factors cannot explain the way that we distinguish the five senses.

1. It is sometimes suggested that, in some sense of know, we all know a simple theory of perception, and this knowledge is supposed to explain our possession of

some practical or conceptual ability. Different writers appeal to such knowledge in order to explain our possession of different abilities. It is not clear, then, that there is a single conception of a simple theory of perception in the literature. Nonetheless, we can ask, and attempt to answer, what the relation is between our concepts of the five senses, and the various kinds of theories which get called simple theories of perception.

A simple theory of perception is simple at least partly because its formulation does not require anything other than a common-sense grasp of psychological and physical concepts: it is a theory that is or could be commonly known, and so if formulated, is formulable in a common-sense vocabulary. It is therefore unlike any scientific theory of perception that we might attempt to formulate. Such scientific theories are not generally known and can only be formulated using a specialised vocabulary and concepts.

Whenever someone claims that we know a theory, we can ask two questions about that theory: first we can ask what is the content of the theory, what is it a theory of; second, how is that content represented, what is the form of the theory and what does knowledge of the theory consist in. Different theorists appeal to a simple theory of perception in different ways and so answer these questions differently. They have different commitments about, for example, the kind of knowledge of the theory that we have, about how the theory is represented, about the content of the theory, and so on. There are serious and difficult questions about what knowledge of a simple theory of perception consists in that anyone who appeals to such a theory must answer. But in what follows I shall be concerned almost entirely with questions about the content of a simple theory of perception, with what the theory represents, rather than with questions about our knowledge of the theory and how that content is represented. I want to know what the relation is between our concepts of the five senses and a simple theory of perception. In particular, I want to know whether our grasp of concepts of the five senses can be explained as consisting in our knowledge of a simple theory of perception. In arguing that it cannot, I shall not be appealing to problems in accounting for our knowledge of such a theory, but in its content and truth. There is no true theory, knowledge of which could constitute our grasp of concepts of the senses. Or so I shall argue.

We can begin by asking what is a simple theory of perception, and what does appeal to a simple theory of perception explain? Gareth Evans introduces the idea of a simple theory of perception as part of an explanation of what it is to have a conception of an objective world. He does so in the context of a discussion about what is required for someone to be able to understand the possibility of something they perceive existing unperceived. It is that idea which is, he thinks (following Strawson (1959, Ch.2)), constitutive of our conception of objectivity,

the idea of unperceived existence, or rather the idea of existence now perceived, now unperceived, is not an idea that can stand on its own, stand without any surrounding theory. How is it possible that phenomena *of the very same kind as* those of which [the subject] has experience should occur in the absence of any experience? Such phenomena are evidently *perceptible*; why should they not be perceived? To answer this question, some rudimentary theory, or form of a theory of perception is required (1985b, pp.261-2).

If all that is required of the theory is that it explain why something perceptible is not perceived, the simple theory can be very simple indeed – it needs only to be the idea of a condition whose satisfaction is required in order to perceive something and which can fail to be satisfied. That is, to paraphrase Evans, one can understand the possibility of its Ø-ing unperceived if that in virtue of which ‘Its Ø-ing’ is true is connected with experience by some condition which is sometimes, but not always, satisfied. A proposition such as ‘it’s Ø-ing’ will be understood to entail that, if that condition is satisfied, it may be perceived to be true and, “[i]n the formulation of the condition there lies a theory, or the form of a theory, of perception” (1985b, p.262). What kind of condition is it that is sometimes but not always satisfied and which would explain one’s failure to perceive something? There are various kinds of explanation of why perceptible phenomena may not be perceived: those that rely on spatial notions, such as the observer’s being in the wrong position or having the wrong orientation, or there being something in the way; those that cite deficiencies in the perceiver, such as that she is inattentive, asleep, unreceptive in the proper modality; and there are those that cite the absence of factors in the world which are

causally necessary for perception, as the absence of light may be cited to explain why we cannot see the table (1985b, p.263). There are, then, many conditions whose failure to be satisfied would explain why we don't perceive something perceptible.⁶⁶

Suppose that Evans is right that our conception of an objective world consists in our grasp of 'a theory, or the form of a theory, of perception'. Would our grasp of such a theory be sufficient to explain the distinctions that we make between the five senses? If all a simple theory of perception consists in is whatever is required to understand why something perceptible is not perceived then there is no reason to think that it would. There need not be anything sense specific in our understanding of why it is that we cannot perceive what is perceptible; we can understand that on the basis of understanding some very general condition necessary for perceiving something with any sense, that is sometimes, but not always, satisfied: a spatial condition such as being in the right place, for example, which is necessary in general for perception, no matter with which sense.⁶⁷ One

⁶⁶ In a later discussion of a similar question Evans' (1982, pp.222-3) emphasis is on spatial conditions. This, it seems to me, constitutes a change. In Evans (1985b) the concern is to explain how one can think of something one perceives as existing unperceived, and the question is whether we can explain that without appealing to spatial notions; an argument that no such explanation could be given would "rest upon the idea that only a spatial theory can satisfy the demand that the factor accounting for the presence or absence of perception of perceptible phenomena should be at once a priori connected with propositions about the world, and yet subject to significant empirical control" (1985b, p.268). It may be that only spatial factors can satisfy this demand, but that needs to be shown. In Evans (1982) the discussion of a simple theory of perception is part of an account of what is required for one to be able to think of oneself first-personally. According to Evans, such thought requires knowledge of what it is for the proposition ΠI am δ_i to be true, and "knowledge of what it is for ΠI am δ_i to be true, where δ_i is a fundamental identification of a person (conceived of, therefore, as an element of the objective spatial order) consists in our knowledge of what it is for us to be located at a position in space...this in turn can be regarded as consisting in a practical capacity to locate ourselves in space by means of exactly the patterns of reasoning that [constitute grasp of the little theory which is required for the idea of an objective spatial world]" (1982, p.223). Space is essential to the later account because of its role in individuating objects; in particular the object which is oneself.

⁶⁷ Even if there is not a single condition which applies to all senses, we needn't suppose that one has to grasp five distinct, sense specific, conditions in order to grasp Evans' "form of a theory" of

must be in the right place relative to something if one is to perceive it, no matter with which sense one perceives it; that one is not in the right place relative to something can explain why one doesn't perceive it. It is for this reason that a conception of space is often claimed to be central to our grasp of objectivity.

It is, however, a small step from thinking in this way to thinking in terms of some more substantial theory. Even a notion as general as that of being in the right place might, when we reflect upon it, look like it would ground distinctions between different senses, since one might think that what counts as being in the right place relative to what is perceived will vary along sense specific lines: the right place for seeing something is different to that for feeling it, and so on.

When we reflect on the reasons why we do not perceive something we can provide different kinds of explanation depending on the sense modality involved. So, although what is required for a conception of objectivity need not be sufficient for distinguishing the senses, reflecting on what is required for perception – what is required in order to perceive something – might do so. Several writers have taken a simple theory of perception to be more like a theory which provides us with knowledge of what is required in order to perceive something, and so as more substantial than whatever is required for an explanation of our conception of objectivity, for an explanation of why we *don't* perceive something.

It might be suggested, then, that we think of a simple theory of perception not as whatever is required to provide an explanation of why something is not perceived, but as setting out the conditions that must hold if one is to perceive something. The shift is from giving some explanation of why something is *not* perceived, to explaining what is required in order to perceive something. These aims are similar, but not equivalent – a grasp of the latter is sufficient for a grasp of the former, but not *vice versa*. This is a significant difference; it's the difference between some, perhaps very general, condition that is necessary for perception and a set of conditions which are, other things being equal, sufficient for perception, and so between a condition in whose formulation "there lies a theory, or the form of a

perception. Another initially plausible general condition is spatial proximity: one can only perceive things which are spatially proximal to one. But that is not true generally: we can see stars, which are not proximal to us.

theory, of perception” and something more substantial. Johannes Roessler, for example, says that

To have a simple theory of perception, in Evans’ sense, is to understand that perception is subject to certain causal enabling conditions. For example, and most importantly...‘a thinker must be able to think of his perception of the world as being simultaneously due to his position in the world, and the condition of the world at that position’. Other ingredients in the theory may include further conditions on particular sensory modalities, for example, the idea that seeing an object requires a clear line of sight between the perceiver’s eye and the object. Having a grip on some of the causal factors affecting perceptibility amounts to a theory in the sense that it involves the ability to give causal explanations (1999, p.58).

He goes on to suggest that:

A subject’s grasp of a simple theory of perception has an important bearing on her ability to exercise control over her perceptual experience... [O]nly when interrogative attention is informed by a simple theory of perception will the question of success or failure arise for the subject herself. Thus a simple theorist can exploit her grasp of the causal conditions of perception by engaging in practical reasoning about what to do in order to acquire the kind of information that she is after, for example where to look to find a particular object. Or again, when an attempt to acquire certain information fails, a simple theorist can recognise this, and consider alternative strategies (p.58).

It is not necessarily part of an explanation of our grasp of the concept of an objective world that we should know what to do in order to acquire any particular kind of perceptual information, so this simple theory of perception goes beyond what is required to explain our grasp of objectivity. When a subject engages in practical reasoning about what to do in order to acquire perceptual information she

is thinking about what would be sufficient for her to have a particular kind of experience, and that requires her to satisfy, or attempt to satisfy, all those conditions which are necessary for a certain kind of perceptual experience – the kind that would provide her with the requisite information, or would enable her to perceive what she wants to perceive. If she is to do this she must attempt to satisfy all those conditions necessary for her to obtain the kind of perceptual information that she seeks, and so she must have some knowledge of what those conditions are. Her knowledge of a simple theory of perception is supposed to constitute this knowledge, and this requires a much more substantial theory than Evans'. John Campbell similarly appeals to what must amount to a substantial theory:

Centrally, we think of our perceptions as causally dependent on the way things are around us: what one sees depends on what is there to be seen... Grasp of spatial relations is central in all these types of thinking...[it] 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, one must look in the right direction, and there must be nothing in the way...one's understanding of the conditions of causal interaction is provided by one's reflective understanding of spatial relations... Someone who has these skills may be said to grasp a simple theory of perception and action. I do not mean that he grasps a schematic theory, stated using abstract concepts. Most people do not have an explicit grasp of the very abstract ideas required in a fully general statement of the theory. What I mean by 'grasp of a simple theory of perception and action' is better described as a skill: the ability to generate causal explanations of particular perceptions... (1994, pp.205-8).

If we think of our perceptions as causally dependent on what they are of – on what's there to be perceived – then someone who grasps a simple theory of perception in Campbell's sense knows what they must do – what conditions must be satisfied – in order to causally interact with the world in the way that is required for them to perceive things in it. These conditions will be different, both Campbell and Roessler suggest, for different senses.

If we think of a simple theory of perception as a general statement of conditions that must be satisfied in order to perceive something, then could knowledge of such a theory explain the fact that we distinguish five different senses, and explain what our possession of concepts of those senses consists in? How could a simple theory of perception provide an explanation of that? It seems likely that a general statement of the theory will be such that it distinguishes different kinds of conditions whose satisfaction enables us to perceive something. Grasp of such a theory could explain the distinctions we make between the senses if there are five such kinds of condition corresponding to the five senses that we distinguish. Someone could be said to possess concepts of these senses just in case they apply these concepts to instances of perception or to particular perceptual experiences on the basis of judging someone (themselves included) to be perceiving something in a way that requires the satisfaction of the relevant kind of enabling conditions. This capacity to apply concepts to perceptions or to perceptual experiences could be explained by their knowledge or possession of the simple theory of perception which articulates the different kinds of enabling conditions. So we could explain the way that we distinguish different perceptions or perceptual experiences into the perceptions or experiences of five different senses by appeal to our knowledge of a simple theory of perception. I have described the form a simple theory of perception might take and suggested that knowledge of a suitable theory might explain the way that we distinguish the senses. The question we now need to address is whether in fact it does so.

We could provide this kind of explanation without having to suppose that people have an explicit grasp of a general statement of the theory. It may be that we have a practical capacity to make the distinctions, a capacity which could be articulated and made fully explicit. Anyone who attempted to explain the distinctions we make between the five senses in this way could, however, plausibly

claim that we can at least partially articulate the theory. People generally can say – can explicitly articulate – something about what has to be the case if they are to see something, or to touch something, over and above (I suspect) the ability to generate causal explanations of particular perceptions.⁶⁸ Whatever someone who appealed to a simple theory of perception to explain the distinctions should say about this does not in any case matter from the point of view of my argument: nothing I say depends on claims about how the theory is represented, on what kind of knowledge of the theory people in fact have.

There are many other questions that we might ask about the relation between a simple theory of perception and our concepts of the five senses. For example, if the simple theory of perception is supposed to explain our grasp of concepts of the senses, should we think of the concepts as concepts *of* kinds of enabling conditions, or as concepts *of* the underlying mechanisms which explain the differences in enabling conditions? This question is analogous to that in the philosophy of mind between functionalists who think that mental states are functional properties and those who think that they are whatever instantiates or realises those properties. We don't need to decide this question because, I shall argue, our grasp of concepts of the senses is not constituted by a grasp of a simple theory of perception.^{69,70}

If we are going to explain our grasp of the distinctions between the senses by appeal to a simple theory of perception then two conditions must be satisfied. The first is simply that there must be a true theory, of the relevant kind, which distinguishes five senses. That is, a true statement of the simple theory of perception must be such that it would distinguish five kinds of (disjoint) conditions

⁶⁸ Whether or not they can do so is an empirical question which has not, so far as I know, been tested.

⁶⁹ Someone who wanted to explain our grasp of the concepts of the senses in this way is, of course, obliged to answer this question. The question is also related to one in the philosophy of science about the semantics of theoretical terms. For the debate in the philosophy of mind, see Lewis (1972), Putnan (1967), and Kim (1998).

⁷⁰ Note that if there is no true simple theory of perception which distinguishes the senses, and there is, as I shall argue in the next chapter, no scientific theory which distinguishes them, then

any of whose independent satisfaction would enable the subject to perceive something, and these conditions must be such that they can plausibly be identified with (or associated with) the five senses that we commonly distinguish.

Secondly, if knowledge of the simple theory of perception is going to explain our grasp of the distinctions between the senses then it must be possible to formulate that theory independently of making those distinctions. If that were not possible, then the theory would presuppose the distinctions that appeal to the simple theory of perception is supposed to explain. To put the point differently. If our grasp of the different senses consists in our knowledge of a simple theory of perception then it must be possible for someone to know a simple theory of perception prior to grasping the distinction between the different senses. This is a point about what would constitute an adequate explanation of our knowledge of the distinct senses. It may well be true, for example, that the enabling conditions for seeing something are different to those for touching something, but it doesn't follow from that that a theory of those conditions can explain either the difference between seeing and touching or our knowledge of what that difference consists in. It wouldn't do so if the only thing common to the conditions which enable us to touch something, and in virtue of which they form a set of conditions necessary for touching something, is that they are conditions which enable us to perceive things in the way we call touching. That is, we might only be able to explain what these enabling conditions all have in common, by appealing to the fact that they are the conditions that must be satisfied in order to touch something, rather than explain what it is to touch something in terms of an independently identifiable set of enabling conditions that must be satisfied in order to perceive something. The enabling conditions must fall into groups corresponding to the senses independently of our distinguishing different senses if a simple theory of perception is to explain those distinctions.

So the question that we need to ask is: would a general statement of a simple theory of perception – of the conditions that must be satisfied simply to enable us to perceive things – be such as to distinguish different kinds of enabling conditions

functionalism about the senses will be false; i.e. the view that our concepts of the senses are functional concepts, and that the different experiences involved can be picked out functionally.

corresponding to the five senses? If so, then we would, in effect, be able to give a common sense theory of the five senses: a common sense theoretical account of what the senses are. We can perhaps think of the simple theory of perception as a common sense functional theory of these mechanisms, in as much as our practical application of the theory will involve our making distinctions between the operations or uses of different mechanisms. To that extent, the theory will be a common sense functional theory of the mechanisms of perception.⁷¹ Such a theory need not say anything about what constitutes the different mechanisms – nothing about the psychology or physiology of the senses – over and above what can be judged to be the case in everyday contexts on the basis of observing people.

Remember that an account of the distinction between the senses must explain more than just what makes the perception of particular objects a perception involving a particular sense, it must explain in virtue of what a perception of the property of an object involves a particular sense, and it must explain what all the perceptions of a particular sense have in common in virtue of which they are perceptions of that sense. The suggestion is that the perception of, say, the shape of something is an instance of seeing the shape if the conditions whose satisfaction were sufficient for the perception of the shape on this occasion were those necessary for seeing. If this is going to explain the way we distinguish the senses, then there must be a single set of conditions necessary for seeing something, another for touching something, and so on for the other senses. Of course, if we have already classified perceptions into groups corresponding to the five senses –

⁷¹ Would the existence of five kinds of enabling conditions imply the existence of five distinct underlying sensory mechanisms? One might think that it would. After all, we often individuate sensory mechanisms in functional terms: so if there are five different kinds of causal enabling conditions that produce perceptual experiences that would provide the basis for the functional specification of five different sensory mechanisms. But this is not necessarily true. When we look at how sensory systems are actually instantiated it may turn out that a single mechanism (individuated in some appropriate way) subserves two different kinds of function (individuated in terms of enabling conditions). This is so because it is possible to get functional dissociations without dissociations in the underlying mechanisms which sub-serve those functions (for a discussion, see Shallice 1988, Ch.11). Nonetheless, given the sorts of reason we are likely to have for grouping together enabling conditions into kinds – in terms perhaps of the involvement of different parts of the body, hands, eyes, noses, and so on – it does look like there will be different mechanisms involved.

into visual perceptions, tactual perceptions, and so on – then we will be able to identify a single set of conditions necessary for, say, seeing. It will be that set of conditions which must be satisfied in order to have a visual perception of something. But the question is whether we can explain how we classify perceptions into such groups merely by appealing to the conditions necessary for perceiving. The answer, I shall argue, is that we cannot: we cannot identify all those conditions necessary for, say, touching something, independently of knowing which perceptions are tactual perceptions. Consequently, we cannot explain how we distinguish the senses by appealing to such conditions. We can see this by considering some examples.

We distinguish the sense of touch from the sense of taste, feeling something from tasting it, although the enabling conditions for both kinds of perception are similar: both, for example, require contact with the object perceived. There is a difference in that, in the case of taste, a certain kind of contact is required: the object has to be placed in the mouth or to come into contact with the tongue.⁷² We can therefore appeal to the different enabling conditions to distinguish tasting something from feeling it. But there are differences in the conditions that must be satisfied in order to tactually perceive various different kinds of thing as great as those between the enabling conditions of taste and touch. If that is so, then appealing to such conditions will not explain why we group all tactual perceptions together as tactual. There is nothing in common to the enabling conditions of all those perceptions we classify as tactual which could explain why they count as tactual perceptions.

It will be sufficient for the point that I am making if I describe three kinds of tactual perception for which this is true; there may well be many more, and there are other examples involving different senses. The first is the perception of the temperature of something or some place. We can feel the temperature of something by touching it with our fingers, and there may be circumstances in which one perceives the temperature of something in this way and one perceives nothing

⁷² According to one interpretation of Aristotle, he suggested that we should count taste and touch as a single sense because they both require contact with the thing perceived. See Sorabji (1971, pp.69-73).

else about that thing. One might, for example, very briefly touch the surface of a hot iron. Equally, there are times when one can tactually perceive some object or property of an object, such as its shape, and yet not perceive its temperature. One might be wearing gloves, for example. Usually, we treat the perception of temperature and of shape as both kinds of tactual perception, as both involving the same sense. Yet the conditions that must be satisfied in order to perceive them – the enabling conditions – are different in each case. Why, then, according to the account I am considering, do we regard both kinds of perception as tactual; why don't we distinguish the perception of temperature as involving a distinct sense from touch, as we distinguish the perception of flavour as involving a distinct sense from touch? Merely appealing to the enabling conditions involved will not answer that. This question is even more pressing, I suggest, when we bear in mind the fact that we talk of feeling the temperature of something even when we are not in direct physical contact with it (c.f. Gibson 1966, pp.129-30). We feel the heat of the fire, for example, or feel that the room is cold. We usually treat perceptions such as these as involving the same sense as our perceptions of temperature by touch, yet one of the enabling conditions for feeling temperature by touch is not satisfied.

Secondly, consider what one must do in order to feel the weight of something in contrast to what one must do in order to feel the texture of something. To feel the weight of something one must lift or heft it; to feel the texture of something one must usually run one's fingers over its surface. That one has to do these different things is a consequence of the fact that there are different sensory mechanisms involved in the perception of each kind of feature. Again, different conditions must be satisfied in order to perceive these different features – our perception of them has different enabling conditions. Why, then, do we group instances of our perceptions of these different properties together as perceptions involving a single sense? Merely appealing to the enabling conditions is not sufficient to explain that.

Thirdly, the kinds of features that one can perceive of an object one touches, and the accuracy with which one perceives them, depends on whether that object simply comes into contact with one's skin, or whether one actively touches it with one's hands (or with some other part of one's body: both the mouth and feet

can be used in active touch).⁷³ Active touch involves the grasping of an object, and allows one to perceive a variety of properties imperceptible by passive touch, properties like rigidity, softness, and so on. We are much better, too, at recognising the shapes of things if we actively touch them and explore them with our hands, than if they are simply pressed up against our skin. Again, we can think of these two kinds of touch as involving different kinds of enabling conditions. The conditions that must be satisfied in order to actively touch something are different to those that must be satisfied in order to passively touch something. In virtue of what, then, are they grouped together as conditions whose satisfaction constitutes the sense of touch?

If we have already distinguished the senses, then we can explain why we group these enabling conditions together as the conditions that must be satisfied in order to touch something. But these examples show that a mere grasp of enabling conditions for perceiving something gives us no independent understanding of the way that we distinguish the senses into five. There are differences in the enabling conditions for the different senses – for perceiving things with each of the different senses – and a grasp of a theory of these conditions may explain our knowledge of what we have to do in order to, say, see something, but unless we already have some grasp of the distinction between different senses a grasp of a simple theory of perception will not explain what the enabling conditions of a particular sense have in common.

(Notice that this conclusion does not impugn the use to which Campbell, Roessler, and others want to put a simple theory of perception. It would merely tell against explaining the way that we distinguish the senses by appeal to a simple theory of perception and neither Campbell nor Roessler claim that we can do that. Despite that, the fact that there are different enabling conditions for different senses might encourage us to think that we can appeal to a simple theory of perception of the kind that they describe in order to do so.)

⁷³ For a further discussion of perception by touching see Gibson (1966, Chs. VI and VII) who stresses the role of active touch in tactual perception, and see Lederman and Klatsky (1987, and

2. To return to the two conditions that must be satisfied if we are to explain the senses by appealing to a simple theory of perception. Can we explain why a simple theory of perception distinguishes five different groups of enabling conditions corresponding to each of the five senses? We could do so if enabling conditions were grouped relative to the sense organs. Different conditions might need to be satisfied to enable us to perceive with each of our different kinds of sense organ; if so there would be five different kinds of enabling conditions corresponding to the five different kinds of sense organ that we have. To decide whether or not this is the case, we need to say something about the sense organs; in particular, we need to explain what our grasp of the different kinds of sense organ consists in. If the way we distinguish different sense organs is relative to a simple theory of perception or presupposes the distinctions we make between the senses, then we cannot appeal to the distinctions we make between sense organs to explain why a simple theory of perception distinguishes five kinds of enabling conditions. Whether or not it does so is the question that I address in this section.

One reason that the question by what criteria the senses are to be distinguished from one another has been so seldom discussed directly may be that certain answers have seemed to different people so obvious that they did not need to be supported by argument. One answer which can, pre-reflectively (and perhaps even after some reflection), seem obvious is that the senses can be distinguished by reference to the sense organs. Roxbee-Cox, for example, in his discussion of the senses, says that

One familiar type of answer [to the question by what criteria the senses are to be distinguished from one another] is that certain physiological processes, involving certain parts of the body with which we are familiar, make an instance of perception a case of sight; the functioning of other processes, involving other familiar parts of the body, make it a case of hearing, etc. This may be called the 'Sense Organ' view (1970, p.530).

1990) for a discussion of the specific, but stereotypical, hand movements involved in tactual object recognition.

All those writers who have directly discussed the question have rejected the suggestion that the senses can be distinguished by reference to the sense organs. In what follows I will discuss what reasons they have for doing so, and whether they are good reasons.

When we think of the sense organs, what we think of first are the eyes, the ears, the nose, and so on. So we might think that we divide perception into senses because different parts of the body are involved in perceiving. The perception of something, on this view, counts as a case of, say, seeing that thing, in virtue of being produced by the exercise of a particular kind of sense organ, in this case by the eyes, and it's because we separate out the eyes from the ears that we distinguish seeing from hearing. If the eyes, the ears, the nose, and so on were not distinguishable parts of the body we would never have distinguished different senses. The problem with this suggestion is that it doesn't explain why we distinguish the five senses as we do. Although it may be true that the reason we distinguish different senses is that different – perhaps spatially distinct – parts of the body are affected, pointing out that different parts of the body are involved in perceiving is not alone sufficient to explain why we distinguish *five* senses. For there are lots of distinguishable parts of the body – there are hands, and feet, and legs, and torsos, and so on – which are (or, at least, can be) involved in perceiving and yet which we don't take to be distinct kinds of sense organ (Nelkin 1990, p.151, makes a similar point).

It might be suggested that we should take the sense organs to be solely those parts of the body that we can use, or at least control, in various kinds of perception. The problem of the precise identification of the sense organs would then be avoided because our knowledge of the organs of sight will be manifested partly in our ability to use certain parts of our body in perceiving things. Referring to the sense organs would merely be referring, by its most easily distinguished element, to the part of the body that we use for a certain kind of perception. But, again, unless we appeal to the kinds of perception their use produces, this doesn't explain why we should group the parts of the body we can control into five kinds of sense organ.

If appealing to the sense organs is going to explain why we distinguish the senses, then we must be able to explain why we distinguish five kinds of sense

organ. We must, that is, explain why we divide up parts of the body into sense organs as we do. Furthermore, we must provide this explanation in a way which doesn't presuppose the distinctions between the five senses. The danger is clearly this: our aim is to explain why we distinguish five senses, and the suggestion is that we can do so by appealing to the fact that we grasp a theory which distinguishes enabling conditions corresponding to the different sense organs. For this to work we have, in turn, to explain why we distinguish five kinds of sense organ; and yet the explanation of that, one might think, is just what we were initially looking for, and so any appeal to the sense organs is otiose. This is the reason that all those writers who have discussed the question have rejected the sense organ view: appealing to the fact that we distinguish five kinds of sense organ does not, in the end, pull any explanatory weight.

That conclusion might seem to have been drawn too swiftly. It's true that we need to say what sense organs are – what kinds of thing are we talking about when we talk about the sense organs – but why think that we cannot do so? Roxbee-Cox suggests that we can think of the sense organs as more than just distinguishable parts of the body, more than just the eyes, ears, nose, and so on. We can extend this list of body parts to include the physiological mechanisms whose functioning is necessary for the eyes, ears, nose, and so on to function in perception (1970, p.531). The senses, according to one physiological textbook, “are the bodily mechanisms for getting up-to-date information” about one's own body and the surrounding environment (Barlow and Mollon 1982, p.1 and p.10). We might appeal to physiology to tell us what these mechanisms are. If we did this, how could we distinguish the sense organs? We could think of the sense organs as those parts of the body which embody or are somehow related to parts of the physiological mechanisms which perform the function of getting up-to-date information about our surrounding environment. In fact, Roxbee-Cox suggests, were we to do this there would be two ways in which we might distinguish the sense organs from one another. The first would be to take the eyes, or the organs of sight, to refer to the physiological system identifiable by the fact that it includes the eyes; the ears to that physiological system identifiable by the fact that it includes the ears; and so on. The second would be to take the eyes to refer to a physiological

system distinguished from others by some such character as being sensitive to a certain kind of stimulus (1970, p.531).

People may be rather vague about the physical and physiological facts involved in the functioning of the senses and so, since we are attempting to characterise our common-sense understanding of the nature of the sense organs, the first of these two ways seems a more plausible account on this physiological interpretation. There are, nonetheless, good reasons for preferring the second way. Unless the sense organs are thought of as things sensitive to certain kinds of stimuli, the sense organ account will have difficulty in distinguishing between seeing something with the use of the eyes and feeling something with the eyes. That is, we can see with our eyes but sometimes we can also feel with them: we can feel when they are touched by something, and we can feel if there is something – an eyelash or piece of grit – in them. If we think of the eyes as including whatever that system is which is identifiable by the fact that it includes the eyes, and of seeing as perceiving with the eyes, then feeling something in one's eye would count as seeing it. Picking out a sense organ in this way would not allow us to make a distinction that we can in fact make, that between seeing something and feeling something in one's eye, and therefore cannot explain how we make that distinction.⁷⁴ Perhaps we shouldn't take this objection too seriously since, whilst it may be true that we can perceive our eyes – feel something in them and scratching them – that is not the same as perceiving something with the eye. When something presses up against one's eye, or some grit falls into and scratches it, one doesn't really perceive the thing touching it or the grit scratching it, rather one perceives one's eye – one perceives a certain part of the body as being scratched or irritated. If we think of the senses as essentially telling us about our environment, rather than about our body, then we can reply to the objection in this way and we can explain why we don't count feeling something in our eye as seeing, even if we think of the sense organs as picked out in the first of the ways Roxbee-Cox describes.

There is, however, a different way of understanding the objection. If we take the eyes, or the organs of sight, to refer to the physiological system identifiable

⁷⁴ This objection is due to Roxbee-Cox (1970, p.53), but has been made by others (see, for example Leon 1988, p.255).

by the fact that it includes the eyes, then the fact that the eyes can both detect light and detect contact suggests that they are part of two different physiological systems: one which detects touch, the other light. Attempting to pick out the eyes as a kind of sense organ in this way will, therefore, necessarily fail: it will not pick out a single or unique physiological process, and so not a single kind of sense organ.

Whether or not that objection is suasive, it seems right that, when saying what the sense organs are, we should appeal to their function rather than to any facts about the actual physical or anatomical structure of whatever it is that realises that function. There are two different kinds of reason for this. The first concerns how physiological kinds are individuated, the second more specifically about the senses.

How do we decide whether two sense organs are of the same physiological kind? Our decision about whether two particular sense organs – perhaps those belonging to two different species of creature – are of the same kind is not determined by whether they are of the same physical kind. We ought to allow, it might be claimed, for the possibility of creatures with sense organs exhibiting a very different physical or anatomical structure to our own, which we would have no hesitation in claiming to be of the same kind as ours (c.f. Leon 1988, p.255). Whether two sense organs are in fact of the same kind cannot, however, be decided simply by appealing to facts about what we would have no hesitation in saying about the senses of other creatures; we must appeal to our scientific theories of physiology, which may be revisionary of the judgements that we have no hesitation in making.⁷⁵ Such theories may abstract away from the actual physical structure of a sense organ, and describe it in more functional terms.

We might accept that our judgements about the senses of other species could be mistaken, and so the fact that we judge creatures to have eyes independently of knowing whether or not they have the same kind of physiology as we do doesn't, in itself, tell against a physiological interpretation of the sense organs and the corresponding explanation of the way that we distinguish the senses. There may, nonetheless, be reasons for rejecting that account. We would allow that

⁷⁵ We might, for example, not hesitate to say that flies have eyes, and yet we could be wrong: physiology might tell us that they don't, or don't in the sense that matters to physiology.

someone who had a damaged eye replaced with a fully functional prosthetic device could still see, irrespective of whether physiology would classify that device as of the same kind as a real eye. If so, it is not the fact that some physiological system is connected with what is actually an eye that makes a particular perception an instance of seeing, nor the fact that it is connected to what physiology tells us is an eye. At most it could be the fact that it is connected to something that functions in a suitably similar way to the way our eyes function. Our judgements about what kind of replacements of the sense organs could be made consistently with maintaining their normal perceptual function are therefore independent of questions about whether such replacements are of the same physiological kind as the original sense organ. Therefore, if such judgements are correct then whatever it is that determines that a particular perception is an instance of seeing, it cannot be that it is produced by a particular kind of sense organ, physiologically individuated. Either we must give up the suggestion that the senses can be distinguished by reference to the sense organs, or we must find some alternative way of understanding the sense organs.

Considerations like these lead Mark Leon to draw the conclusion that,

if we mean by 'eye' that physiological mechanism which we possess, then eyes are not necessary for sight; any prosthetic or functionally equivalent device will suffice. But if we mean by 'eye' some device which subserves a certain function, then eyes will be necessary for sight, but we then have no way of specifying when an item is an eye without appealing to something other than the nature or structure of that item. We would have had to have distinguished the senses prior to determining whether an item was an eye or not (1988, p.255).

Roxbee-Cox objects to the physiological interpretation of the sense organ view for similar reasons. He claims that although scientists might be interested in classifying the physiological processes involved in perception, and so might make distinctions between the senses on the basis of their physiology, there is no reason to think that a concern with classifying the physiological processes involved underlies the use and value of the distinctions as they are ordinarily made. If that is true, then what, he asks, "is the principle for classing together certain parts of the body as organs of

sight, other parts as organs of hearing, etc?" (1970, p.532). In particular, "[w]hy is the familiar classification not arbitrary, whereas a classification that grouped together the left eye, the left hand and the left ear as the organs of a certain kind of perception would seem arbitrary and absurd?" (p.532).

There is, he thinks, something right and non-arbitrary about the way we in fact group the sense organs, and any explanation of that will have to appeal to something other than the sense organs themselves. If we have to explain why we distinguish the sense organs in a way which corresponds to the five senses by appeal to something other than the sense organs, then we cannot appeal to the sense organs in order to explain why we distinguish the senses as we do. Whatever it is that explains the way we distinguish and classify sense organs will explain the way that we distinguish and classify the senses, and any appeal to the sense organs will be explanatorily otiose. Roxbee-Cox is surely right that a classification that grouped together the left eye, the left hand and the left ear as the organs of a certain kind of perception is arbitrary and absurd. It's hard to imagine ever grouping those together as a single sense. But, why does it seem unimaginable or absurd? It is a condition of any adequate account of the distinction that we make between the senses that it should be able to explain why such a classification does seem absurd. According to Roxbee-Cox the familiar classification of the sense organs is not arbitrary because:

there is some element common to the members of the groups of parts of the body that we call organs of sight, feeling, etc., that is not shared by members of the group of 'left-hand organs'. To find such a feature, we shall...have to fall back on such considerations as the character of the experience or the properties perceived, that are associated with the functioning of the various organs (p.533).

The suggestion is that, because there is nothing about the sense organs as we ordinarily think of them that would enable us to explain why the left-hand grouping is absurd, an explanation of that will have to appeal to differences in the kinds of perceptual experiences we have or in the kinds of properties we perceive. The reason we think the left-hand grouping of the sense organs is absurd is that in fact

we pick out the sense organs relative to the kinds of perceptual experiences – visual, tactual, auditory, and so on – that they function to produce. We specify a kind of sense organ – the eyes, say – as a kind of thing which functions to produce a visual experience. If we individuate sense organs by their function in this way, then the sense organs themselves drop out of the explanatory picture: they are just whatever perform the relevant function, where the relevant function can be specified independently of knowledge of what actually realises that function. Some groupings of the sense organs – the left-hand grouping included – seem absurd because they cross-cut our prior classifications of functions.

To summarise: we can only appeal to the sense organs to distinguish the senses if we can explain why we distinguish the sense organs as we do. The mere fact that they are distinguishable parts of the body is not enough; appealing to physiological accounts of the sense organs doesn't seem to work either. Two arguments have been put forward against such accounts: firstly, we don't have a common-sense interest in the physiology of the senses and so we cannot explain the fact that we make common-sense distinctions between the senses as the result of such an interest; and, secondly, our judgements about which instances of the perception of something are instances of, say, seeing that thing, do not correspond with the classifications of perceptions we would make if we made them on the basis of the physiology of the senses. So whatever the basis of our classification is, it's not physiological.

I ended the last section by asking whether a simple theory of perception would distinguish five kinds of enabling conditions corresponding to the five senses, and suggested that it may do so if there exist five kinds of sense organ each with different kinds of enabling condition. The argument of this section has attempted to establish that we cannot explain why we distinguish and classify the our sense organs as we do independently of an explanation of why we distinguish and classify the five senses as we do. We cannot, therefore, explain our grasp of the distinction between the senses by appealing to a simple theory of perception which distinguishes different kinds of enabling conditions corresponding to each of the senses. We cannot do so because the distinctions we make between the sense organs can only be explained by appealing to our grasp of a theory – a simple theory – of how they function in perception to produce experiences corresponding to the

five senses. If that is right, then we cannot appeal to the sense organs in order to explain why a simple theory of perception distinguishes five kinds of enabling condition.

In this chapter I have been discussing whether our grasp of the way that we distinguish five senses can be explained by appealing to our grasp of a simple theory of perception, and I have argued that it cannot because there is no true theory which could explain what the differences between the five senses consists in. It may be thought that this is because the senses are distinct physiological systems or distinct psychological capacities whose nature can only be revealed by scientific investigation. Whether or not they are is the question that I consider in the next chapter.

7. The Senses as Natural Kinds

It is not implausible to think that each of our five senses is constituted by a distinct kind of brain mechanism which realises our capacity to perceive in that particular way; we might think, then, that the senses are to be distinguished by reference to the internal mechanisms which constitute them. What makes it the case that one is perceiving with a particular sense is that a particular kind of mechanism is involved: a particular kind of mechanism produces one's perceptual experience. Thus when one apparently sees something, one's experience is produced by the visual mechanism; that it is produced by that mechanism is what makes it a case of apparent seeing; and a judgement to the effect that one sees something is a judgement that one is having a perceptual experience produced by such a mechanism.

There is, however, a *prima facie* objection to this suggestion. If there are internal mechanisms which constitute our five senses, then their existence and nature will be determined empirically: knowledge of them must consist in knowledge of some theory describing how the brain functions. But that is something that most people will not know; in fact, it is likely that, at the present time, no one knows exactly what mechanisms are involved in perception. In that case, how could concepts of such mechanisms be part of our everyday thought and talk? Concepts of the senses *are* part of our everyday thought and talk, so it looks as though, even if there are different mechanisms underlying our capacity to perceive, such mechanisms play no role in our everyday thought and talk about the senses. If one lacks a concept of, say, the visual mechanism, then a judgement that one sees something cannot be a judgement that one is having an experience produced by the visual mechanism.

This conclusion is premature. It has become common to claim that someone can have the concept of a natural kind – a natural kind concept – in advance of knowing the correct scientific account of the nature of that kind. So, for example, someone can have the concept of gold without knowing exactly what gold is, without knowing what determines whether something is actually gold. In

general, then, that someone is ignorant of the nature of something doesn't mean that they cannot make judgements involving concepts of that thing (Kripke 1980, and Putnam 1975). We can, therefore, defend the claim that our concepts of the senses are concepts of different kinds of brain mechanisms if we can show that they are natural kind concepts.

In this chapter I describe what conditions are necessary for our concepts of the five senses to be a natural kind concepts, namely, that there actually be natural kinds corresponding to them. If the senses do form distinct natural kinds, then a true psychological explanation our capacity to perceive will distinguish five sub-capacities corresponding to each of the senses. In order to decide that question, we need to look to a true and complete psychological theory of perception to determine whether it distinguishes five kinds of sensory capacity. Unfortunately, nobody has yet produced such a theory. Nonetheless, I suggest, if we know what form such a theory would take, we can attempt to determine whether the evidence that we do have about how the mechanisms of perception function supports the claim that there are five such capacities. I describe two different models of psychological explanation and set out the conditions that would have to be met by an explanation of perception if it were to support the claim that our concepts of the senses are natural kind concepts. I then describe some evidence which suggests these conditions are not met and hence that there are not five kinds of sensory mechanism, and hence that the our concepts of the senses cannot be natural kind concepts.

1. A natural kind concept applies to something in virtue of its possession of certain fundamental properties which determine whether that thing is a member of the kind in question. There are a number of ways in which someone could come to possess a concept of a natural kind. Someone could learn of the fundamental properties constitutive of that kind by, for example, learning a particular scientific theory; or it could be that ordinary language already embodies a theory – a kind of folk theory – of the kind in question, so that in learning what a word means one learns about the kind. But there are many cases in which it seems plausible that someone can possess a natural kind concept in advance of knowing what determines membership of the kind. My use of the term 'natural kind concept' will be restricted to these

kinds of cases, rather than, as it is sometimes used, to include any concept which is a concept of what is in fact a natural kind.

According to one well known model, the terms of ordinary language can be correlated with the natural kinds discovered by science because they involve an indexical component of meaning. When we use a natural kind term we refer to whatever natural kind the paradigmatic instances of the extension of the term belong to. These paradigm instances are identified either ostensively or operationally by appealing to a conventional idea, of what something looks like or of what it is, that we associate with the term. The kind is defined as those individuals that bear the appropriate "sameness relation" to the individual paradigmatically identified in this way. An ostensive definition, for example, of a kind concept might go like this,

Suppose I point to a glass of water and say 'this liquid is called water'. My ostensive definition of water has the following empirical presupposition: that the body of liquid that I am pointing to bears a certain sameness relation to most of the stuff I and other speakers in my linguistic community have on other occasions called 'water' (Putnam 1975, p.225).

Our use of a natural kind term commits us to there being a kind to which the thing we refer belongs – it must stand in a same kind relation to other things. This commitment may not be met, but when it is met we can successfully refer to a kind without knowing what it is that determines something as a member of that kind. In saying that a concept is a natural kind concept we are committed to there being same kind relations between instances of the kind, so for this to be the correct account of a particular concept there needs to be sameness relations of the appropriate sort between the individuals the concept supposedly picks out.

This model of natural kind concepts is due to Putnam and Kripke (see Kripke (1980), and Putnam (1975)).⁷⁶ According to Putnam there are four components to the meaning of a natural kind term: what he calls a syntactic marker, a semantic marker, a stereotype, and an extension. So the term 'tiger', for example, might have as syntactic marker "noun", as semantic marker "animal", as stereotype "four legged, yellow coloured animal with stripes, and so on", and an extension determined by the fundamental nature (usually micro-structural or other theoretically discovered feature) of tigers. Putnam's distinction between stereotype and extension reflects a distinction between competence in the use of a term and full knowledge of the meaning of a term. Someone can be competent in the use of a term even if they are ignorant of its extension for as long as they know the first three components associated with that term. Putnam explains a stereotype as those features that must be known by any competent speaker of the language, irrespective of whether it provides a good guide to the actual extension of a term. A stereotype is a conventional idea of what something looks like or of what it is, and "the central features of the stereotype generally are criteria – features which in normal situations constitute ways of recognising if a thing belongs to the kind" (p.249, and p.230). We can use terms of whose extension we are ignorant because of what Putnam calls "a division of linguistic labour" (p.227). If we need to know that terms apply to correct kinds then we must appeal to experts who know the fundamental properties of the kind in question and are able to determine whether something really is a member of the kind. We can never be sure, of course, that the experts really do know the meaning of the term, because there is no guarantee that they know the fundamental properties of the kind in question.

Putnam's account gives us a way of tying ordinary language classifications of things to scientific classifications of things into kinds, or to those that would eventually be provided by science. There are various objections that have been made against his account, and so against any view which appeals to it in order to explain our possession of particular concepts. For present purposes I am simply

⁷⁶ Putnam and Kripke present their accounts as accounts of natural kind terms or words, rather than concepts. In most of this chapter I shall talk in terms of concepts; in what follows nothing turns on the distinction between concepts and terms.

going to assume these objections could be met and that Putnam's account, or something like it, is correct. My objections to the claim that our concepts of the senses are natural kind concepts will not depend on the details of any particular account of the semantics of natural kind concepts; my point is simply that it's not, *prima facie*, implausible to suppose that our concepts of the senses could be natural kind concepts.

How, then, can we decide whether our concepts of the senses are natural kind concepts? One might object to the attempt to identify our concepts of the different senses with concepts of scientifically discoverable kinds because one is, in general, sceptical of the existence of any scientifically discoverable natural kinds. One might, for example, be sceptical of the idea of a thing's having certain fundamental properties which determine it as the kind of thing it is (Mellor 1977); I am not going to argue in that way. Alternatively one might claim that, although within a certain area of classification the relevant natural kinds exist, the classifications of ordinary language do not coincide with those of a scientific taxonomy. This would be the case if our everyday classifications of things cross-classify with the scientific classifications, or make distinctions where there are none, or fail to distinguish kinds distinguished by science (Dupré 1993). Exactly this, I shall argue, is true of our concepts of the senses: it is for that reason that we would be wrong to think that they are natural kind concepts.

For an identification of everyday concepts with concepts of natural kinds, in particular for an identification of our concepts of the senses with concepts of natural kinds, to be plausible, we have at least to show that there are, or could be, the relevant natural kinds corresponding to our concepts of different senses. That means, I suggest, that there must be five kinds of sensory mechanism each of which realise a different sense. How do we determine whether there are five kinds of sensory mechanism? We can look to a true psychological explanation of our capacity to perceive; such an explanation can be viewed as providing a psychological theory of perception, and so as telling us what the capacity to perceive is. If this explanation distinguishes five sensory mechanisms, then a necessary condition for the identification will have been met. But, more importantly, if such an explanation does not distinguish five kinds of mechanism then it will not be plausible to identify

our concepts of the senses with natural kind concepts and we must look elsewhere for an account of the nature of the five senses.⁷⁷

2. We postulate psychological mechanisms in order to explain the psychological capacities of an organism, like the capacity of an organism to perceive its environment. What form should the explanation of a psychological capacity have, and how in general can we explain such a capacity? How it is that the postulation of a mechanism can explain a psychological capacity?

It is widely supposed that we can explain psychological capacities *functionally*. The account of functional explanation worked out in most detail is due to Cummins and Fodor (See especially Cummins (1983) and Fodor (1968) and see Fodor (1983)). In what follows I begin by giving a brief account of functional explanation, and then go on to assess what kind of evidence there could be that the senses form five functional kinds.

According to Cummins, we can explain a complex *psychological* capacity in the same way that we explain any other complex capacity: by analysing it into simpler elements. There are two ways in which this can be done. We can analyse the capacity itself, give what he calls a *functional analysis* of the capacity; and we can analyse the thing or, to use Cummins's expression, the system, which has the capacity: that is, we can give a *compositional* or *systems analysis* of the system which has the capacity. Often an adequate explanation will require both kinds of analysis since what makes an explanation the explanation of the capacity of a particular system is that the system actually realises that capacity. This sets an empirical constraint on the analysis of any capacity of a system: the analysing capacities must be shown to be capacities of the system which has the capacity. This point needs emphasising; it has various significant consequences.

Systems analysis explains how a particular system works by decomposing it into parts, it "explains S's possession of [a property or capacity] by appeal to the properties of S's components and their mode of organisation" (Cummins 1983,

⁷⁷ Fodor suggests that we think of "the kind predicates of a science [as] the ones whose terms are the bound variables in its proper laws. I am inclined to say this even in my present state of ignorance,

p.15). So when we give a compositional or systems analysis of some thing – some system – we explain why it has the particular capacity it does by analysing it into components or parts, and then explaining its possession of the capacity in question by appealing to properties of the components and their organisation. Analysing a system in this way has an explanatory value “when we come to see that something having the kinds of components specified, organised in the way specified, is bound to have the target property,” namely the capacity that we want to explain (1983, p.17). This process of analysis is recursive: since the components we use to analyse a complex capacity will often themselves have capacities or properties which we want to explain, we can analyse them in turn, and we can analyse the components of these components, and so on.

Suppose that we want to explain the capacity of a mechanical watch to indicate the time by the position and movement of its hands. We can analyse the watch into its various parts – its spring, cogs, escapement, and so on – and then show how, given their organisation, the interaction of the watch’s parts makes its hands move so as to indicate different numbers on the watch’s face corresponding to the time of day. We will have explained the capacity of the watch to indicate the time when we have shown how that capacity results from the properties and organisation of the various parts of the watch’s mechanism. Such an analysis would show how the watch works, but we might go on to explain the various properties of its parts which enable them to play the role they do in the watch’s mechanism – the rigidity of the cogs, the flexibility of the spring, and so on – and to do so we would apply the same kind of analysis to each part of the watch that we applied to the whole watch. (Of course, to provide this further explanation would require a different kind of knowledge: we would have to know about the materials which make up the parts of the watch, and why they have the properties they do.)

As well as analysing a system into its component parts we can analyse the system’s capacity. This kind of analysis Cummins calls functional analysis (p.28). The functional analysis of a capacity consists in analysing it into a number of simpler or less problematic capacities in such a way that the organised activity of the

accepting the consequence that it makes the murky notion of a kind viciously dependent on the equally murky notions of law and theory. There is no firm footing here.” (1974, p.132).

analysing capacities amounts to the activity of the analysed capacity. The explanatory value of functional analysis consists in the fact that it helps us to understand the operation of a complex capacity by breaking it down into a series of relatively simple steps, so that we can see how a series of relatively simple capacities, which may be of a different kind to the capacities we are attempting to explain, can together produce very complicated ones.

This kind of functional analysis of a capacity is often a preliminary step to explaining how a system possesses the complex capacity in question. Such an explanation works by showing how the simple analysing capacities described by the functional analysis are themselves realised by relatively unsophisticated parts of the system which possesses the capacity. So we begin the explanation of a complex capacity by analysing it into a number of simpler capacities, and then explain how some system realises or possesses the complex capacity by showing that various component parts of the system themselves realise or possess the simple capacities described by our analysis. Thus functional analysis goes together with systems analysis when we show that the analysing capacities are capacities of components of the system. Fodor suggests that psychological explanations employ just this methodology; such explanations, he says,

have characteristically exhibited two phases that, although they may be simultaneous in point of history, are nevertheless distinguishable in point of logic...in the first phase of psychological explanation, the primary concern is with determining the functional character of the states and processes involved in the etiology of behaviour...The second phase...has to do with the specification of those biochemical systems that do, in fact, exhibit the functional characteristics enumerated by the phase-one theories. (Fodor 1968, pp.107-9.)

Although often a functional analysis of a capacity goes together with a compositional analysis of the system which has the capacity, with the analysing capacities being capacities of components of the system, the functional analysis of a capacity need not have this form. Sometimes we will analyse the capacity of a system into other capacities which are capacities of the system as a whole, and not

capacities of any of its components. To see this, consider, for example, the capacity of a cook to bake a cake. Such a capacity can be analysed into a (sequence of) simpler capacities – to break eggs, to follow instructions, to mix ingredients together, and so on. But these simpler, analysing, capacities are not capacities of some part of the cook, they are just capacities that the cook has; capacities, we might say, of the whole cook. So to provide an explanation of the capacity of the cook to bake a cake, we would have to provide a functional analysis of the capacity, but not a compositional analysis of the cook. (Compare this with a production line which produces cakes: different parts of the production line are responsible for different parts of the process, so in order to understand the capacity of the production line to produce cakes we would need to analyse the capacity into capacities of parts of the production line). We might, of course, go on to explain the cook's capacity to break eggs by appealing to capacities of parts of the cook.

We can provide a functional analysis of a capacity without reference to an instantiating system since, in effect, all we do is re-describe the capacity; consequently "functional analysis puts very indirect constraints on compositional analysis" (Cummins 1983, p.29). It is important, however, to keep in mind the difference between a functional analysis which analyses the capacity itself into simpler *component* capacities, and an analysis which analyses the capacity into simpler capacities which are *not* component capacities. The distinction is important when we are attempting to explain the possession of the capacity by a particular system, the human brain for example.

In practice, when we attempt to explain the capacity of a system we often need to analyse the complex capacity of the system into simpler capacities of the system as a whole before attempting any compositional analysis of these capacities. Indeed, it will often be possible to analyse the complex capacity of a system in different ways and, since the different analyses will have different implications for the structure of the system which instantiates them, it will be important to distinguish these analyses when we come to explain the instantiation of that complex capacity by a particular system. The same capacity might, for example, be the product of two distinct and simpler capacities of the system as a whole, or it might be a single complex capacity of the system – we need to know which before attempting to provide any further compositional analysis.

A more concrete example of this is provided by a class of distributed networks which are commonly used for simulations of cognition. Such networks are set up so as to have a particular steady state function – to produce a certain kind of output given a certain input – but the network can also ‘learn’ to produce a particular output given a particular input. We can analyse the complex capacity of this system into two simpler capacities – to produce a steady output and to learn – which are both capacities of the whole system or network, not of components of it. Yet the same complex capacity could be implemented by a system having two components each with a simpler – distinct – capacity (see Shallice 1988, p252).

If it is to be explanatory, the functional analysis of a capacity of a particular system or thing must eventually terminate in capacities whose instantiation can be explained by an analysis of the system itself. We must be able, in the end, to explain how that particular system has the capacity. Since there will usually be more than one way in which a capacity can be analysed, what makes an analysis an explanation of the capacity of a particular system is that the system instantiates the analysing capacities. If the system doesn’t have the capacities that analyse the capacity in question then we will not have explained how this particular system has the capacity. In order to substantiate the claim that we have provided an analysis of some capacity of, say, the human brain, we need to show that our analysis is in fact instantiated by the brain. If, for example, we analysed some complex capacity into two simpler, and distinct, capacities of the whole system working in parallel, we would look for evidence that these capacities are in fact independently instantiated by the system, and our claim to have provided an analysis will be undermined if we cannot show the analysing capacities to be independently instantiated. This requirement is what sets an empirical constraint on correct analysis.

This method of explanation can be used to explain any complex capacity and it can, in particular, be used to explain the possession of a capacity by some system. We can think of the psychological capacities of an organism as an instance of the possession of a complex capacity by a system, and so we can use this combination of functional and systems analysis to explain the psychological capacities of an organism. A theory which explains the psychological capacity of an organism like, for example, the capacity of an organism to perceive its environment, will first analyse the capacity into simpler capacities whose joint operation constitute

the capacity to perceive. Such a theory might analyse the capacity to perceive into distinct capacities to see, to hear, a proprioceptive capacity, and so on. Since the goal of an analysis of psychological capacities is to explain how those capacities are instantiated by a particular organism the theory could go on analysing these capacities by breaking them down into various component capacities, each of which could then be shown to be instantiated by parts of the system. In the case of psychological capacities, this point will have been reached when the capacities can be shown to be instantiated by parts of the nervous system and brain of the organism.

A complete functional explanation of some capacity of a system has the form of a hierarchy of levels. The top level is simply a description of the capacity to be explained; an analysis of this capacity constitutes the next level down, and each of the lower levels in turn provides an analysis of the level above. Each analysis of a higher level capacity is constrained by the requirement that the components of the analysis actually be instantiated by the system. If we cannot show that the lower level description is a description of capacities of components of the system then we have a reason to reject it as an explanation.

How could we decide whether there were five kinds of sensory mechanism corresponding to the five senses? If a complete and adequate explanation of the capacity of humans to perceive their environment analysed the general capacity to perceive into five distinct kinds of capacity instantiated by parts of the nervous system and brain then we might identify these capacities and their underlying mechanisms with the five senses. Such an analysis would vindicate the claim that there exist five kinds of sensory mechanism which together constitute our capacity to perceive and so could form the basis of an argument to the effect that our concepts of the senses are natural kind concepts. Unfortunately, we don't have this a complete and adequate explanation of the human capacity to perceive nor, given the current state of neuropsychology, are we close to getting one; so we cannot determine whether or not there are five kinds of sensory mechanism in this way.

We do, however, know what form such an explanation would have to take if it were to analyse the capacity to perceive into five distinct sensory capacities and, given that the capacities of an empirically adequate analysis must be instantiated, we know what kind of structures or mechanisms there would have to be in the brain

were such an analysis correct. That means that, in advance of having complete explanation of perception, we can look for evidence either that the brain does or does not have the relevant kind of structures or mechanisms. Such evidence would constitute evidence for or against the existence of five kinds of sensory mechanism independently of having a fully worked out explanatory analysis of the capacities in question; independently, that is, of having a complete explanation of the human capacity to perceive. Evidence of this kind might never be such that we could say for sure, in advance of a complete explanation, that there were five kinds of sensory capacity; but it could, I suggest, provide convincing evidence that there were no such capacities.

If our concepts of the five senses are concepts of natural psychological kinds, then the brain must actually instantiate five kinds of sensory capacity. The existence of more or fewer than five kinds of capacity would undermine such an identification.⁷⁸ An explanation of our capacity to perceive would therefore have to analyse that capacity into five distinct and simpler capacities, and it must show that these capacities are instantiated by distinct parts of the brain. So, for example, it must show that the brain instantiates a capacity to see – a visual capacity. Our everyday concept of seeing is the concept of a single sense, so this concept could only be identified with the concept of a kind of psychological capacity if the brain instantiates a single visual capacity – if there is, in the brain, a single visual mechanism. It follows that any explanation of capacity to see must itself be such that it analyses it only into simpler *component* capacities; were it to analyse the capacity to see into two or more independent capacities whose combined parallel operation constitutes the capacity to see then it would not be plausible to claim that seeing is the exercise of a single kind of capacity. Rather, there would be two kinds of capacity, both of which we might call visual capacities, but neither of which would be *the* capacity to see. The existence of two such visual capacities would undermine the identification of our everyday concepts of the senses with concepts of kinds of psychological capacity. We can look for evidence that there is a single

⁷⁸ Actually, that's not quite right: if there were more than five such capacities we could still perhaps claim that our concepts of the senses were concepts of natural kinds if we could provide some further explanation of why we only have concepts of five of them.

visual capacity instantiated by the brain, so here we have a clear example of how we might actually go about deciding whether our concepts of the senses are natural kind concepts.

What form would this evidence take? What would constitute evidence that the brain did instantiate five sensory capacities? What, in particular, would constitute evidence against the existence of a single visual capacity? A capacity is individuated by its functional role, its function being to map types of input onto types of output in a certain way.⁷⁹ If a capacity is an input-output mapping, then to claim that a sensory modality such as vision is a single kind of capacity instantiated by the brain is to claim that there is some mapping of input states onto output states of the brain which instantiates it. We will have identified such a capacity when we have identified its inputs and its outputs and the relation between them; two capacities will be identical only if they map the same kind of inputs onto the same kind of outputs in the same way. If the correct analysis of our capacity to perceive analyses it into five distinct capacities, one of which is a capacity to see, and further analyses that capacity as a single kind of capacity which is not further analysable except into component capacities, then we should expect to find it instantiated in the brain as a single kind of mapping of inputs onto outputs. If, on the other hand, our capacity to see is analysed onto two distinct parallel capacities, then we should expect to find it instantiated as two distinct input-output mappings in the brain.

In general if the explanation of our capacity to perceive analyses it into five distinct kinds of sensory capacity – five input-output mappings – then we should expect to find five such capacities instantiated in the brain as distinct mappings of input states onto output states. Were we to do so then it might be plausible to claim that our concepts of the senses are concepts of these capacities: that our concepts of the senses are concepts of these kinds of psychological capacity.

Evidence for or against the claim that the senses are natural kinds will therefore take the form of evidence that the brain does or does not instantiate

⁷⁹ Cummins says that “to ascribe a function to something is to ascribe a capacity to it that is singled out by its role in an analysis of some capacity of a containing system”, and that “X has a disposition to d if X would manifest d were a certain range of events to occur. To explain a disposition is to explain why d comes about when precipitating conditions occur”. A disposition is the same as a capacity for Cummins, see p.195 n.1.

distinct sensory capacities, or distinct kinds of input-output mappings. If vision is a single capacity then it must consist of some single mapping of inputs onto outputs. If there exists no single kind of input, or no unique kind of output, then vision would not be a single kind of capacity. Evidence against such uniqueness would be evidence against the claim that vision is a single kind of capacity, and if it can be shown that vision is not a single kind of capacity then that would be sufficient to show that the senses in general are not natural kinds.

Marr, in his discussion of the function of vision (a function whose workings he goes on to describe in detail) takes the input to the visual capacity to be fairly obvious:

A process may be thought of as a mapping of one representation onto another, and in the case of human vision, the initial representation is in no doubt – it consists of arrays of image intensity values as detected by the photoreceptors of the retina (1982, p.32).⁸⁰

The input to the visual capacity is, then, the pattern of light that falls on the retina. That seems right: in order to see something we must detect light emitted by or reflected from it, and that light is detected when it falls on the retina. So whatever the output of vision may be, the input must be the light that is detected by the retina. We might think, too, that something similar is true of the other senses; it is fairly easy to distinguish different sensory ‘transducers’ – those parts of the body – the sense organs – which are sensitive to various different kinds of stimuli and hence detect different kinds of information about the physical world. All the information we have about the world comes to us thanks to the operation of these transducers, so we know that the input to whatever perceptual capacities we have can be no more than what is detected by them; and it might seem safe to infer from this that the input of each of the different sensory capacities is just what is detected by each sense organ – in the case of vision, arrays of light intensity values.

⁸⁰ Marr talks of processes where I talk of capacities; the difference is not important from the point of view of my argument.

But we want to know whether there is a *single* capacity corresponding to each transducer, and that depends in part on how the transducer itself functions. We know, for example, that the retina contains neurones which are sensitive to different features of the visual array, and we know too – a fact “that is often not appreciated” (Milner and Goodale 1995, p.3) – that projections from the neurones in the retina travel to a number of different targets in the brain.⁸¹ So, that the retinal image is the input to a single capacity, rather than that different properties of the image are inputs to distinct capacities, is an assumption which may turn out to be false. It is possible that we could find evidence that there are several kinds of visual output, which would be evidence that there is not a single input-output mapping for vision. That would suggest that vision is not a single capacity.

3. Suppose that we find evidence that there is not a single input-output mapping for vision. That would seem to be enough to undermine the claim that there is a single visual capacity, and so undermine any identification of our concept of vision with the concept of a natural kind. The question is not, however, quite so clear cut. I have been supposing, as do many philosophers, that psychological explanation is a form of functional explanation and so have described what kind of evidence would show that there could not be a functional explanation of perception which distinguishes five kinds of sensory capacity. But whether or not we can explain the capacity of some system functionally is not itself a trivial question: it depends, in part, on the kind of structure the system has. An adequate functional explanation is one whose analysing capacities can be shown to be instantiated by components of the system. But suppose the organisation of the system whose capacities we are attempting to explain doesn't have the kind of structure that allows us to distinguish it into components. In that case it will never be possible to show that a functional analysis is instantiated by the system – there simply would not be components of the system to instantiate the relevant capacities. Given that there is this empirical assumption build into functional analysis, evidence that some system doesn't

⁸¹ Physiological and anatomical studies have distinguished several distinct classes of retinal ganglion cells each of which appears to be involved in the analysis of a different aspect of the visual scene.

instantiate a particular capacity is not necessarily evidence that the system lacks that capacity: it may show instead that the system lacks the kind of structure that makes functional explanation appropriate.

According to the account of psychological explanation that I have been discussing – functional explanation – we can explain a psychological capacity by analysing it into sub-capacities; and we can show that the analysis is the correct analysis of the capacity of a particular thing or system by showing that the analysing capacities are capacities of parts or components of the system, that they are instantiated by components of the system. Suppose, for example, that we analyse the complex capacity of some system into two parallel capacities each of which are capacities of the system as a whole. To show that this analysis is correct we would need to show that both capacities are in fact instantiated by the system in the way described by our analysis. That means the system must be shown to instantiate two independent capacities (together with any component capacities we postulate in further analysing these capacities). If two capacities of the system are independent of one another then, since component capacities are, in turn, individuated in terms of their role in the analysis, each set of their component capacities must themselves be independent of one another. So in order to show that such an analysis is the correct analysis of the capacity of the particular system that we are attempting to explain we need to show that the system has components which instantiate those capacities in the way described by the analysis. If the analysing capacities are not instantiated then we cannot claim to have provided an analysis of the capacity of this particular system: although we might have provided an analysis of the capacity, and so given an explanation of how some system could have this capacity, we would not have explained how or why this particular system actually has the capacity. If we can't show that the system really instantiates two independent capacities, then our analysis will not be the correct explanation of the capacity of this system. Exactly this kind of procedure is described by Fodor:

The cells form two channels of information which remain partially segregated through higher cortical regions. See Cowey (1979) and Logothetis (1990).

The image that suggests itself to many psychologists is that of opening a “black box”: having arrived at a phase-one theory of the kinds of operations performed by the mechanisms that are causally responsible for behaviour, one then “looks inside” to see whether or not the nervous system does in fact contain parts capable of performing the alleged functions...The physiological psychologist’s task of determining what, if any, organisation into subsystems the nervous system of an organism exhibits is precisely the problem of determining whether the nervous system has subsystems whose functional characteristics correspond with those required by antecedently plausible psychological theories (1968, pp.109-110).

If this model of the functional explanation of a psychological capacity is to be applicable to the various psychological capacities – capacities like vision and memory – that we ordinarily attribute to people, then the system which instantiates these capacities – the human nervous system and brain – must have a certain kind of structure, it must have what Fodor calls a “modular” structure: that is, its various distinct capacities must be implemented by collections of sub-components – parts of the nervous system and brain – which are themselves independent of one another.⁸² People do often think of the brain in this way. It is common, for example, to view the brain as analogous to a machine, to a computer, say, and yet it is an entirely non-trivial fact about machines that they have distinct component parts.⁸³

Were the brain to lack this kind of structure then we could not explain its capacities functionally because there would not be any empirical way to determine which of two alternative analyses of the same brain capacity was correct. Although

⁸² For Fodor a module is actually a sub-component with a special set of properties: “A module is, inter alia, an informationally encapsulated computational system – an inference making mechanism whose access to background information is constrained by general features of cognitive architecture, hence relatively rigidly and relatively permanently constrained.” (1990, pp.200-1).

⁸³ Fodor compares the nervous system to a computer and, elsewhere, to an engine. They differ in their levels of complexity, but not in the kind of structure they have (1968, pp.109-110). Cummins’s (1983) examples are similar: an electronic circuit, a computer, and so on.

we could provide functional analyses of its capacities, we would not be able to show that such analyses were in fact analyses of capacities of the brain; we could not do so because we could not show that the capacities were actually instantiated either by the brain, or by its component parts, for the simple reason that it lacks such parts.

Given the empirical commitments of this form of psychological explanation and in advance of a complete explanation of perception, evidence against the existence of five sensory capacities corresponding to the five senses will be equivocal. On the one hand, we could take it to be evidence that the correct functional explanation of the senses will not, in fact, distinguish five sensory capacities; or, on the other hand, we could take it as evidence that the human nervous system and brain lacks the kind of structure required if this model of psychological explanation is to apply to it.

Functional explanation sets a very strong constraint on correct analysis. In practice the constraint seems too strong. Couldn't there be other, less demanding ways of empirically determining which of two (or more) alternative analyses of a psychological capacity is a correct analysis of the capacity of the brain. In practise psychologists can and do construct explanations of people's psychological capacities, and they do so by constructing models of how those capacities are realised in the brain. Such explanations often distinguish capacities at one level even when they are produced by the operations of sub-capacities which are not entirely distinct at a lower level. If that's so then, in practice at least, psychologists don't necessarily individuate analysing capacities by their role in the analysis of higher level capacities, since they may also play a role in the analysis of other capacities that we distinguish at the higher level. In effect we can see functional explanation as placing too strong a requirement on an adequate explanation – a requirement that makes it too easy to find evidence against a proposed explanation of a capacity. There's reason to weaken the instantiation requirement on an adequate explanation, and hence to raise the level of evidence required to show that a particular analysis is not instantiated by a system. What form do these explanations take, and what kind of evidence do psychologists in practise appeal to in providing an analysis of a particular psychological capacity? The best way to answer these questions is to consider a concrete example. I've simplified it somewhat, to make the point more clearly.

We may not know exactly how to characterise human memory – common sense tells us that it stores certain kinds of information, but not what or when or how. So if we want to explain memory we must begin by attempting to characterise the capacity to remember in detail. On the face of it, there is more than one way to think about a person’s capacity to remember things. We might think of it as a single, general, capacity which can be used to remember any kind of information; or we might think of it as consisting of several distinct capacities to remember different kinds of information: the capacity to remember the way to get into the centre of town may be a different capacity to the capacity to remember how to perform a mathematical calculation, or to remember the date of one’s birthday, or to remember the smell of a flower. And so an explanation of memory must begin with the best, most detailed, description that we can give of the capacity to remember that most people actually have: we need to describe what people can actually remember, and what they cannot remember. It may be far from obvious what people can in fact remember; as Churchland comments, what the mind is doing “even described at the level of input-output functions of the system – is not an observational matter, to be read off simply by looking at the behaving organism. Rather, it is a deeply theoretical matter. Some initial theory is essential to get the whole enterprise going...”⁸⁴ The theory here will take the form of hypotheses about the role of memory in various kinds of behaviour. That is, we hypothesise psychological capacities which are responsible for people’s behaviour. The capacities we hypothesise must be sufficiently complex to account for whatever behavioural abilities people are shown to possess. Producing this kind of description is likely to involve a certain amount of empirical investigation: we need to experiment, to test people, and so on, in order to determine what they can do. Producing this description is what Fodor calls phase-one of psychological explanation: we construct a psychological theory of memory. Once we have analysed the capacities people have to remember things we attempt to locate the mechanisms which instantiate these capacities. This is Fodor’s phase-two:

⁸⁴ Churchland (1989, p.374) and see Fodor (1968, Ch.3).

the physiological psychologists task of determining what, if any, organisation into subsystems the nervous system of an organism exhibits is precisely the problem of determining whether the nervous system has subsystems whose functional characteristics correspond with those required by antecedently plausible psychological theories...it is clear that a psychological theory that attributes to an organism a state of process that the organism has no physiological process capable of realising is ipso facto incorrect...If no such mechanisms exist, then the [analysis of that capacity] is the wrong model for the functional organisation [of that capacity] (1968, pp.109-110)

I have argued that our failure to find such mechanisms does not show that the analysis of the particular capacity we are attempting to explain is incorrect: it may simply be that the capacity is instantiated by a system which lacks the required kind of structure. How, then, do psychologists determine whether an analysis – a theory of memory – is correct?

When they begin constructing an account of people's capacity to remember – a psychological theory or model of memory – psychologists don't just use evidence of what people can do, they use evidence of what they can't do: their theories are often based on studies of the abilities of people with brain damage. By studying what abilities remain intact in the absence of others – by looking at how abilities dissociate – when the brain is damaged it is possible to learn about the structure of the intact capacity. Shallice provides a succinct explanation of the kind of methodology employed:

The importance of dissociations stems from an inferential asymmetry between associations and dissociations, if observed impairments faithfully reflect damage to an underlying modular system. If one patient shows an association between two types of deficit and a second shows a dissociation, with one of the abilities being preserved, then a simple explanation of the overall pattern exists. The observed dissociation can be presumed to arise from a lesion that has affected

only one side of a functional line of cleavage in the modular system; the association is presumed to result from a lesion that has crossed this line.⁸⁵

In fact it is more important to look for double, rather than single, dissociations between two abilities. Two abilities are doubly dissociated if each one can be impaired without the other being so. This suggests that different underlying psychological capacities or mechanisms are required for the two abilities. A double dissociation is more significant than a single dissociation because a single dissociation is compatible with the possibility that the same capacities underlie the two abilities, but that the impaired ability simply taxes those capacities more heavily (and so stops working first, or works much less well, when they are damaged).

This approach has proved useful for understanding the mechanisms of memory.⁸⁶ In some people with damaged brains, memory impairment – amnesia – occurs as a circumscribed disorder, and without any cognitive impairment. Although we tend to think of memory as a single capacity to remember things, the study of people with amnesia has provided evidence for distinguishing two kinds of memory: short and long term memory. People with amnesia can remember things for short periods of time, but not for longer periods. One psychologist who has studied memory concludes that such results “suggest a distinction between at least two kinds of memory” (Squire 1989, p.504) – or two kinds of capacity: the capacity to remember things for short periods of time, and the capacity to remember things for longer periods of time. People with amnesia are in fact often able to learn things – they are often able to learn how to perform certain motor, perceptual, and cognitive tasks. They can, in other words, still remember some kinds of things for longer periods of time, which suggests that we should distinguish different kinds of long term memory. It is this the kind of evidence psychologists appeal to in attempting to characterise a particular psychological capacity that they want to

⁸⁵ Shallice (1988, p.35). This methodological approach is relatively recent and interest in it has greatly increased over the past 25 years or so. Shallice's book is an excellent discussion of its theoretical underpinnings.

⁸⁶ Some of the relevant evidence here is summarised in Squire (1989).

explain; it's evidence that, in the case of memory, appears to show that what we thought of as a unitary capacity is in fact the joint operation of distinct capacities.

Although this evidence *suggests* that we need distinguish different kinds of memory it is possible that the kind of fragmentation of capacities revealed by amnesia has no functional significance and does not reflect the underlying structure of the mechanisms of memory. In order to show that human memory does have this kind of structure and that the analysis is correct we need to show that these different capacities are in fact instantiated by the brain. If we are not able to do so, then we would have to go back and revise our initial characterisation. Typically psychologists will look to anatomical evidence to confirm an analysis; as Squire comments, "to understand how the brain actually accomplishes learning and memory, it is important to identify the specific brain structures [involved]". We might not be able to identify such structures, in which case the anatomical studies would undermine our initial distinctions and show that they have no real significance. In order to show that our distinction between two kinds of memory was correct we would, according to the model of functional explanation, have to identify two distinct independent capacities instantiated by the brain, but that is too strong a requirement.

The kind of evidence – evidence of dissociation – which leads psychologists to characterise memory as two distinct capacities can be produced by the interaction of systems of component capacities which are not themselves independent of one another. That is, our analysis of a capacity of some system, based on evidence from dissociations, might distinguish distinct capacities of that system whose sets of analysing capacities – in this case, capacities of various physiological structures in the brain – are not entirely independent of one another. Two sub-capacities might interact and yet still be sub-capacities of different and distinct higher-level capacities. This would not be possible according to the model of functional analysis. According to that model, such interaction between two component capacities would be sufficient to show that they are components of a single capacity, and hence that any analysis which counted them as components of distinct capacities would be wrong. In practise, psychologists do not take evidence of lower level interaction of this sort to necessarily rule out a higher level distinction.

There are various kinds of brain structure which could produce the kind of dissociation which would count as evidence of distinct capacities of the system.⁸⁷ Suppose, for example, that each of two sub-capacities has a different input and output, and that each can function effectively without the other, but that there is some interaction between them. The interaction might function in such a way that, for example, the two capacities are not able to produce conflicting outputs. There could be significant interactions between two such sub-components, and strong connections between them. According to the model of functional explanation, such interactions would rule out viewing these sub-capacities as components of distinct higher-level capacities. And yet in the situation in which one of these sub-capacities is damaged, the other may continue to function; so it may be appropriate to treat each sub-capacity as a component in two different higher-level capacities. Whether or not it is appropriate to treat these sub-capacities as distinct from one another may depend on whether we can specify the functions of each of the capacities of which they are part independently of one another. That is, evidence from a higher level may lead us to treat components that interact at a lower level as distinct. Alternatively, suppose that two capacities are realised in overlapping areas of the brain. One capacity A might require regions X and Z of the cortex, and another capacity B might require regions Y and Z. We might count X and Y as contributing to the function of two independent capacities, even though they are not realised independently of one another. We might still regard A and B as distinct.

These kinds of structure (and there may be many others) are examples of systems which are not modular in the way required by functional analysis, and yet that the components are not independent is not taken to undermine the claim that the capacities they realise are in fact distinct. Saying this does not imply that the lower-level structure and organisation of a system places no constraints on an analysis of the capacities of that system. We can still say under what circumstances we should count capacities as genuinely distinct or not, and there can still be cases where the absence of a certain kind of lower level organisation or structure would rule out some putative analysis of a system's capacities.

⁸⁷ For further discussion, see Shallice (1988, Ch.11).

We can view the relation between a single sub-capacity and the rest of the system in which it is embedded as on a continuum. In providing an analysis of some psychological capacity our decision as to whether to treat two capacities as genuinely distinct or not will be determined by the extent to which the operation of a sub-capacity depends on outputs of other sub-capacities of the system of which it is a part, and how they relate to the state of the rest of the system. There may be certain relations between sub-capacities which are of far greater importance than others, so that we can group together and distinguish components on the basis of the strength or importance of the connections between them. For as long as all the components of a capacity have strong connections to one another, and weak connections to the components of other capacities, we can view the components as genuinely instantiating that capacity and the capacity as genuinely distinct from others.

We can say, then, under what circumstances some proposed analysis of a capacity has a genuine functional significance that reflects the underlying structure of the mechanisms of the capacity without having to suppose that the system whose capacity we are attempting to explain has a modular structure.

The different systems I have described all allow that component capacities have differing degrees of functional specialisation. Component capacities can be specialised even when they do not have a unique function. If the brain is not organised in a modular way, but does have capacities, and systems of capacities, that are functionally specialised, then we can explain the capacities of the brain, like vision, by analysing them into systems of functionally specialised components or capacities, and we can group components or capacities according to their degree of functional specialisation: we might, for example, call a component, or a capacity, a component of the visual system because it makes a functionally significant contribution to the visual system, hence to the visual capacity. In the case of memory, what would show that the characterisation of memory as two independent capacities was correct would be the discovery of two functionally specialised memory systems – one for short term memory, one for long term memory.

The difference between this approach – explanation in terms of functional specialisation – and functional explanation is that this approach allows that there is some explanatory value to an analysis which shows how complex capacities can

emerge from the interaction of functionally specialised components, even when we cannot precisely characterise the function of those components and the contribution they make to the system as a whole, and even when they don't stand in the very tight relationships to the analysed capacity required by functional analysis.

This approach still does make empirical assumptions, namely that the brain is has functionally specialised regions. But there is ample evidence that the brain is so structured. The very fact that specific impairments are produced by damage to specific areas of the brain supports it.

4. What would be evidence that there are not five kinds of sensory capacity? Initially I described how the requirement of a functional explanation of perception that capacities which are distinct at a higher level must be instantiated by distinct sub-capacities at a lower level could provide such evidence, since evidence of interaction between capacities is evidence that they are not component capacities of distinct higher level capacities.

But if we reject that requirement then what form could the evidence take? We should look for evidence that there are five distinct sensory capacities, evidence in the form of dissociations. And we should ensure that these capacities are functionally significant, that they are realised in functionally specialised parts of the brain. Evidence that there is not a single input-output mapping would not be sufficient to show that there was not a single capacity. The descriptions of functionally specialised structures above include examples of structures which interact: that a capacity has more than one input or more than one output doesn't imply that it is not a single capacity since only some inputs and outputs need be functionally significant. I will end by describing some evidence which undermines the claim that vision is a single sensory capacity and, a fortiori, the claim that our concept of the senses are natural kind concepts.

Why do we tend to think that vision is a single capacity, that we have a single visual ability? That we do is, we might think, indisputable. There is, as I described earlier, what appears to be a single input to vision; whether vision is a single capacity will depend on whether this input is the input to more than one functionally specialised capacity. Evidence of more than one functionally significant

visual output would be evidence of more than one specialised capacity, and there is evidence of more than one functionally significant output.

We commonly suppose that the function of vision is to provide some sort of internal representation of the world around us which can serve as a perceptual foundation for thought about, or action on, various aspects of the world. It is this kind of picture of the function of vision which supports the idea that there is a single visual capacity. There is, however, considerable empirical evidence that this picture is mistaken. There is a great deal of evidence, for example, that the visual system is highly modular; that is, that distinct functions are computed by distinct, functionally specialised, systems. Different modules are responsible for processing information about colour, motion, pattern, form, depth, and various other attributes of the perceived world (Zeki 1993). The function of many of these sub-systems is, as yet, only poorly understood; some have functionally significant inputs from more than one sensory modality (those sub-systems responsible for speech perception are one such example. They have inputs from both audition and vision. That makes their status as part of either a visual perceptual capacity or an auditory perceptual capacity problematic).⁸⁸ What reason is there to think that all these sub-systems produce outputs which make a functionally significant contribution to a single visual representation of the world? If they do not do so, then it would be wrong to describe vision as a single perceptual capacity, rather than as a collection of independent and functionally specialised capacities. In fact, there is evidence that vision involves at least two independent functionally specialised capacities, sometimes (inaccurately) described as capacities for computing 'what' something is and for computing 'where' something is, capacities which are realised by two independent processing streams in the brain.⁸⁹ That these capacities are functionally independent of one another is suggested by certain characteristic patterns of dissociation which I describe in some detail below.

⁸⁸ See Massaro (1998). For a discussion of multi-modal perceptual systems see Stein and Meredith (1993, esp. Ch.II).

⁸⁹ For the original distinction see Mishkin, et al (1983); for a recent survey of some relevant evidence, see McCarthy (1993). Most of the evidence that I describe below is drawn from Milner and Goodale (1995).

Consider our ability to pick up something that we can see. To pick up an object successfully the hand must be correctly oriented with respect to the object, the size of the grip must be adjusted appropriately, and the fingers must be placed at geometrically appropriate opposition points on the object's surface; furthermore, if the object has a particular use, it must be picked up in a way appropriate for its use: a spoon must be grasped by its handle and not its bowl. When we attempt to pick up an object we not only reach towards its spatial location, but we also shape our hands and fingers in anticipation of the object's shape, size, and orientation.

When we grasp an object we don't wait to touch it before shaping our grasp, instead the construction of a motor programme directed to an object uses visual information about the object's properties to preshape our grasp. The grip that is formed by the hand when it is in contact with the object is actually the end product of a sequence of movements which start when the hand first starts to move towards the object; our grasp is preshaped on the basis of visual information.

Preshaping involves a progressive opening of the grip, with a straightening of the fingers, followed by a closure of the grip until it matches the object's size. The opening of the hand must be matched to the size of the goal object. In a normal grasping movement the hand begins to open almost as soon as it begins to move towards the object, and reaches maximum aperture about two-thirds of the way through the action. This maximum, which is larger than the object itself, is strongly correlated with the size of the object (Jeannerod 1986). That this pattern of finger movement which arises prior to, and during, grasping reflects visual information from the object is demonstrated by the fact that the amplitude of the grip aperture during grip formation co-varies with visually perceived object size (Jeannerod 1997, p.35).⁹⁰

In the normal case, then, visual information determines various aspects of an object directed action. This is perhaps what we would expect: it's in agreement with common sense intuitions to say, for example, that we reach in the direction that we do because that is where we see the object we want to pick up to be. Thus

⁹⁰ Several other aspects of reaching and grasping are also visually determined including, interestingly, the initial grip force. See Gordon (1991).

we might take these empirical studies as confirmation of our common sense intuitions about the connections between perception and action.

That there are connections between perception and action is indisputable. We think that we can explain properties of a persons action by appeal to the nature of their perceptual experience; so we take it that a person knows how to act with respect to something that they perceive *because* they perceive it. We know, for example, the direction in which to move to pick up an object because we have visually based knowledge of the location of the object.

Peacocke, in discussing these connections, suggests that when asked to direct the beam of a spotlight onto a particular tree he can see an “explicit statement of a subject’s practical reasoning” would be of the following form:

The subject forms an intention with the content (1).

(1) I will move my arm in the direction of that tree.

He also knows from his perceptual experience (2).

(2) That tree is in direction *d* (identified egocentrically...[on the basis of perceptual experience]).

So he forms the intention with content (3).

(3) I will move my arm in direction *d*.

He can then carry out this intention without any further practical reasoning. (Peacocke 1992, p.94).

It is part of this common sense picture that the subject could have had the very same perceptually based knowledge of his environment without acting in any way – because he lacked the appropriate intentions, say – and that he could have used this very same knowledge for some other purpose, identifying ‘that direction’ as the direction in which he must walk in order to get home, for example. This picture provides us with one reason, then, for supposing that there is a single capacity involved in vision – there is just one functionally significant output which is put to various uses. It is used to make visually based judgements, to act on the basis of what is seen, and so on. That the picture may not be true in general is suggested by the following two cases.

An example of one of the many visual size illusions to which we are subject is provided by the so-called 'Titchener circles' in which two circles of equal size are presented, each surrounded by an array: one of circles smaller than the target, one of circles larger than the target. Typically, people see the target circle which is surrounded by the smaller circles as being larger than the target circle which is surrounded by the larger circles. It is possible to make the two target circles appear to be the same size if their real size is adjusted.

A version of this illusion was used by Aglioti and his colleagues to test which aspects of object-directed action are affected by the illusion (1995). They set up an experiment in which two thin discs were used as the target circles. Subjects were instructed to pick up to disc on their left if the two discs appeared equal in size, and the one on their right if they appeared different in size. The aperture of their grip was then measured during the grasping movement.

During the experiment, the size of the two target discs was varied randomly so that sometimes although the discs were actually different in size, they appeared equal, and at other times although they appeared different they were actually the same size. All subjects were apparently affected by the visual illusion and treated discs which were actually different as perceptually the same, and they treated discs which were actually the same as perceptually different.

The scaling of their grip, however, proved to be immune to the illusion. Even when they perceived the two different discs as identical, they continued to scale their grip appropriately, opening their hand significantly wider for the slightly larger of the two discs. The same actually correct grip scaling was observed when the targets were the same size although perceived to be different. The subjects' grip size was determined by the true size of the target disc, and not its apparent size. "Thus, the very act by which the subjects indicated their susceptibility to the visual illusion (that is, by picking up one of the target circles) was itself uninfluenced by that illusion." (Milner and Goodale 1995, p.168).

This seems to be inconsistent with the common sense picture I described since, according to that picture, the various features of our object-directed actions are determined by our visually based knowledge of the object. Yet here we apparently have a case in which our object-directed actions are correct even though our visually based knowledge of the object is mistaken; so the action, although

visually guided, is not guided by our visual knowledge: the two have come apart. It looks as though there are two functionally independent visual outputs: one used to control actions, the other as the basis of judgements.

A number of studies have shown that subjects are able to update or amend the direction or grip of a grasping movement towards an object if the object is seen to move or change suddenly during the movement. Surprisingly, however, such amendments can be made even when the subject does not perceive the change in the object which is responsible for the amendment.

When we reach towards a target which appears suddenly in the periphery of our visual field we both move our arm, and we move our eyes and head so that the target is foveated. This saccadic eye movement towards the object is usually completed before the arm movement is completed. Goodale and others performed an experiment which showed that information available from a target after it has been foveated can be used by a subject to correct the trajectory of their hand as it moves towards the target (1986).

In the experiment, subjects were asked to move their finger from a central target to a new target which appeared suddenly in their peripheral visual field. In half of the trials, the target stayed in position until the subject had completed the movement. In the rest of the trials, the target was displaced suddenly to a new position. This sudden movement of the target took place during the saccadic eye movement. The two kinds of trial occurred in random order and the subjects were not told that the target would sometimes change position.

In the trials in which the target moved, the final position of the subject's finger was displaced by an amount equivalent to the target displacement: subjects corrected the trajectory of their aiming movement. Despite this, subjects at no time perceived the target as having moved, nor did they detect anything different between normal and displaced target trials. Even when they were told that during some reaches the object would move, and were asked to indicate which ones, their performance was no better than chance.

This result again seems inconsistent with the picture I described. Here we apparently have an example of a visually-guided object-directed action which can be appropriately amended in response to visual information even when the subject is

unaware that any amendment needs to be made. So, once again, the action, although visually guided, is apparently not guided by our visual knowledge.

That these examples are not superficial phenomena, but reflect something about the underlying psychological processes is supported by clinical reports of patients with certain kinds of brain damage. Milner and Goodale describe a patient of theirs, D.F., who has severe difficulties at various perceptual tasks without corresponding deficits in her visuomotor abilities (1995, Ch.5).

D.F. has severe visual form agnosia and is unable to recognise visually presented objects, although she can still detect flashes of light, and fine gratings, though not their orientation, and she has well preserved colour discrimination. Her ability to discriminate between or recognise even simple geometric forms is grossly impaired and she performs very poorly on the Efron test, a task which requires the subject to discriminate rectangles of different proportions (but of the same surface area).

Despite these deficits, D.F. can reach out and grasp objects accurately, and she is good at catching a ball that is thrown at her. She can easily negotiate objects in her path, and she can follow a moving light with her eyes. Although she is poor at reporting the visual qualities of objects – their size and orientation – she is much better at using these qualities to guide her actions. So, for example, in one test a slot was cut in a vertically mounted disc. On different trials the orientation of the slot was randomly set at 0°, 45°, 90°, or 135°. D.F.'s attempts to make a perceptual report of the orientation of the slot showed little relationship to its actual orientation (whether the reports were made verbally or by manually setting the orientation of a comparison slot. Her deficit isn't one of verbal report). Yet when she was asked to insert her hand or a post a hand held card into the slot from a starting position an arms length away, she showed no difficulty, moving her hand or the card towards the slot in the correct orientation and inserting it quite accurately. Video recordings of her movement show that her hand began to rotate into the correct orientation as soon as it left the start position, as is the case for normal subjects.

In D.F. we have someone for whom the visual information available for controlling action and that available for making perceptual reports and for recognising objects has become dissociated by brain damage. What is interesting

about this case is that it does not seem possible to explain what has happened within the framework of a single visual capacity. Information about the same features of visually perceived objects is apparently available to some higher-level processes, but not others.

Yet if there is a single functionally significant output of the visual process, information about particular features will either be present, and so available to higher-level processes, or else absent, and so unavailable. This case, and the examples described earlier, seem to be inconsistent with the claim that there is a single functionally significant output to vision. Instead there must be two or more outputs, one of which is used by the higher-level processes which control action, and the other used by the processes involved in object recognition. But if there are two functionally significant outputs then there are two functionally specialised capacities. Information about shape and orientation is missing from one of these outputs, but not the other – presumably because the damage to D.F.'s brain has affected one capacity but not the other. If this explanation of the dissociation of abilities seen in D.F. is right, then it is wrong to think of vision as a single capacity. And if that is right, then our everyday concept of vision cannot be a natural kind concept because vision itself is not a single kind of psychological capacity.

I have described, in some detail, evidence which suggests that vision is not a single capacity; that we can distinguish at least two independent functionally specialised visual sub-systems. This is not the only evidence to which I could have appealed in support of my claim that vision is not a natural kind. Face recognition provides another example. There is evidence that our capacity for visual face recognition is sub-served by a functionally specialised and independent system; and that it employs distinct and specialised computational processes (see Farah (1995) for a summary of some relevant evidence). This evidence takes the form of patterns of impairment of visual face recognition – a deficit known as prosopagnosia. Such impairments occur without impairments to the ability of people to recognise other kinds of object or otherwise to visually perceive things, suggesting that “prosopagnosics have lost the specific system that is necessary for recognising faces...” (Farah 1995, p.104). That tends to support my claim that vision is not a single kind of capacity: our perception of some of things that we classify as visually perceived depends of the operation of a functionally specialised capacity,

independent of other capacities that we regard as visual; if that is true, then those things that we classify as visually perceived are not produced by a single capacity, hence vision is not a natural kind. I could equally have described other examples, such a speech perception, a process which involves both auditory and visual inputs, and hence whose status as either and auditory or a visual process is problematic (see Massaro 1998); I take the examples that I have described sufficient to establish my claim that vision is not a natural kind.

Why, it might be asked, if vision – and the senses in general – are not unitary sensory capacities, do psychologists and others continue to talk of visual processes, of tactual processes, of auditory processes, and so on, as they do in textbooks and dictionaries? Surely their doing so indicates that they regard these as different kinds of process corresponding to each of the senses. We needn't suppose that it does. Most psychologists of perception are not concerned with explaining the distinctions that we make between the five senses; they are usually concerned to explain our capacity to perceive some particular kind of thing or feature, such as movement, or faces, or objects. That means that they can draw on the common-sense distinctions that we make between the senses in characterising different processes as visual, and the rest. What they call visual processes are either just those processes that play a role in the perception of things that we would normally regard as visually perceived, or else capacities which process a certain kind of information, derived from the retina, say. That doesn't mean that we can identify a single visual process independently of the way we normally distinguish and classify the senses. That psychologists tend to group processes relative to our common-sense distinctions is, we can suppose, either just a matter of convenience or a way of marking that they have a certain kind of input. Nothing follows from it about the structure of the mechanisms involved, or about the functional organisation of the senses.

Even if it is true that vision, and the five senses generally, are not psychological capacities, that still leaves open the possibility that they are some other natural kind, in particular it leaves open the possibility that they are physiological kinds. In order to decide whether or not they are we can look to accounts of the physiology of the senses. Are there different kinds of brains mechanism – different kinds of physiological mechanism – involved in perception?

Undoubtedly there are, but the same constraints apply, if we are to appeal to such mechanisms in giving an account of the five senses of common-sense, as apply to an appeal to psychological capacities. In particular, there must actually be five kinds of physiological mechanism which could plausibly be identified with our common sense distinctions. Are there, then, five kinds of sensory mechanism?

Given the connection between the structure of the system which instantiates them and the psychological capacities they instantiate (which I have describe above), the existence of five physiological brain mechanisms sub-serving perception would imply the existence of five kinds of psychological mechanism sub-serving perception realised by those mechanisms. We have seen that there are not five kinds of psychological capacity, and that implies that there are not five kinds of physiological mechanism either.

Perhaps, though, the relevant mechanisms are not brain mechanisms, but those parts of the body which are connected to the brain, the various sensory transducers or sense organs which detect various kinds of environmental information. If there were five such transducers then we could perhaps identify the senses with them. I discussed this kind of suggestion above when discussing attempts to distinguish the senses by appeal to the sense organs. Then I postponed the question of whether we could explain the senses by appealing to the sense organs understood as physiological kinds. So we can now ask whether there are five different kinds of sense organ – five kinds of sensory transducers which function to pick up information and relay it to whatever higher-level processes occur in the brain. A great deal more is known of the functional physiology of the body than of the brain, and we can simply turn to a textbook of physiology to answer this question.

The answer is clearly no. Physiologists distinguish many more than five kinds of sensory mechanism. One example will suffice to make the point. We might think of the sense organ of touch as the skin. Certainly we talk of feeling or touching many properties or features of things that we detect using the skin. The skin, however, does not contain a single sensory mechanism, but several: “[t]here are at least 15 functionally and morphologically distinct kinds,” including those

which detect temperature and various mechanically sensitive receptors.⁹¹ We regard things that we perceive as a result of the operation of any or several of these distinct mechanisms as things that we feel, so we cannot simply identify feeling with whatever is produced as the result of the operation of a certain kind of sensory mechanism. Similar points apply to the other senses.

In this chapter I have argued that, whatever the five senses are, they are not natural kinds. That is, our concepts of the five senses are not concepts of certain kinds of sensory mechanism which sub-serve perception. My argument for this has not relied on any objection to the idea that we could commonly have knowledge of such mechanisms, but on the fact that there just are no mechanisms of the appropriate kind. Our concepts of the five senses are not, therefore, natural kind concepts. That does not mean, of course, that there are no sensory mechanisms. Nor does it mean that the our distinctions between the senses are entirely independent of the existence of such mechanisms. But it does mean that our concepts of the senses are not concepts of different kinds of sensory mechanism, and it means that appealing to such mechanisms will not explain the way that we distinguish five different senses.

⁹¹ Iggo (1982, p.369). This paper contains a good general survey of the physiology of cutaneous sensory mechanisms. See also Gibson (1966, pp.109 ff.).

8. The Senses as Conventional

In previous chapters I have argued that various attempts to explain the way we distinguish five different senses are inadequate. In this chapter I sketch my own account of the distinction between the senses, and of what the senses are. My explanation develops from a suggestion of Edward Craig's that it can be illuminating to ask what the point of some concept is, what role that concept plays in our lives. I argue that if we ask that question about our concepts of the senses we can find grounds for an explanation of why we distinguish five senses in the way that we do. My account explains the distinction as being a matter of convention. That means that we can explain the way we distinguish the senses without having to suppose that there really are five kinds of sense independent of the fact that we distinguish them. My account is, if you like, a deflationary account of the senses.

In what follows I begin with a brief account of Craig's suggestion. Then I provide an explanation of the point of our concepts of the senses. Given that explanation, I argue, we can appeal to conventions to explain why we distinguish five senses. I end by comparing my view with the Brute View that I described in chapter 5; my account, I suggest, provides a better explanation of the way we distinguish the senses than the Brute View, and so we have reason to reject the Brute View in favour of mine as the correct account of the senses.

1. The standard approach to questions about the concept of knowledge has been to search for the necessary and sufficient conditions for something being a case of knowledge; conditions which match our intuitive judgements as to whether putative examples are examples of knowledge. Craig suggests that we try an alternative approach, that of asking what point the concept of knowledge has for us. He supposes that there may be concepts for which an account of their point will reveal something of the nature of the thing the concept picks out. The concept of knowledge may be one such; it is, he claims, a very widespread concept:

There seems to be no known language in which sentences using 'know' do not find some comfortable and colloquial equivalent. The

implication is that it answers to some very general needs of human life and thought, and it would surely be interesting to know which and how (Craig 1990, p.2).

Since an understanding of its purpose promises to be at least as interesting as any analysis of the conditions of its application, this suggests an alternative way to approach any attempt to understand the concept of knowledge. To pursue this alternative

[w]e take some prima facie plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions would govern its application (1990, p.2).

This approach is not in general guaranteed to lead anywhere interesting. Every language has a word for water, but the purpose of having that word is just to be able to talk about water, "something which every community has an obvious need to be able to talk about." But no thought about or understanding of that purpose will bring us any closer to understanding water. Why isn't the concept of knowledge just like that? The concept of water is, we suppose, determined by the nature of water and our experience of it; the concept of knowledge, Craig suggests, is the concept of a kind of thing which, unlike water, is not entirely independent of our interests and purposes:

Knowledge is not a given phenomenon, but something that we delineate by operating with a concept which we create in answer to certain needs, or in pursuit of certain ideals (1990, p.3).

If this were true, if it is indeed the case that counting certain states as states of knowledge can only be understood in relation to our practise of conceptualising them as states of knowledge, then an answer to the question about the point and importance of that concept to us, of our purposes in calling certain states

knowledge, can be expected to reveal something about the nature of those states of knowledge.

Exactly this is true, I shall argue, of our concepts of the senses.⁹² We can ask what role the concepts of the five senses play in our lives, and why is it important or useful to us to make distinctions between different senses. In what follows I shall argue that there is a connection between the sense with which someone perceives something and what they perceive; and that we can appeal to this fact in order to provide an explanation of the point of our concepts of the senses. Put briefly, perception is an important source of knowledge. Different ways of perceiving – perceiving with different senses – leads to the acquisition of knowledge of different kinds of thing. So, if we are interested in what knowledge of things other people have or what knowledge they lack, then we will be interested in the different ways that they perceive things. It is, furthermore, useful for us to track or remember which ways people perceived things – our concepts of the senses are concepts of these different ways of perceiving. Which of the different ways that people perceive things we track – which senses we distinguish – is, however, arbitrary, a matter of convention.

2. A natural place to begin any attempt to explain the point of our concepts of the senses is by asking what the consequence is of the fact that we see some things rather than feel them.

What explains whether or not we perceive something? Any explanation of what we perceive will appeal to a variety of different conditions. These conditions will include spatial notions, like being in the right place, whether there are obstructions, which direction one is looking in, what one is touching, and so on. It will appeal to various facts about the perceiver, like whether she is paying attention. And it will appeal to facts in the world which are causally necessary for perception, like the presence of light, the absence of background noise and so on. These conditions will vary according to the particular sensory capacity we use in

⁹² Although his approach is interesting, there are reasons to think that Craig is wrong about knowledge. See, for example, Williams (1978, pp.39-45) and Peacocke (1999, pp.35-6). I am going

perceiving; what we perceive – both which particular things we perceive and which way we perceive those things as being – is, in part, determined by which senses we use. Thus, whether you have your eyes open or keep them shut and simply listen, will affect what you perceive; and that you see something rather than touch it will affect what properties of that thing you are able to perceive. Which way you perceive determines what things you are able to perceive and what you are able to perceive about them. Which way you perceive determines, therefore, what perceptual knowledge you acquire. We can use this fact to explain why it should be of interest to us that someone sees something rather than touches it; to explain, that is, why it should be of interest to us with which sense or senses they perceive.⁹³

We have an interest in what knowledge other people have and what knowledge they lack. What someone does is often determined, at least in part, by what they know or believe; in particular it is determined by what they know or believe about their environment. Therefore, in order to understand, predict, and explain the behaviour of others successfully we must be able to determine, more or less reliably, what it is that they know or what they don't know (often an explanation of what someone did will appeal to the fact that they didn't know or didn't realise that such and such was the case. Someone's failing to see something, for example, can make reasonable a course of behaviour which might strike us as entirely unreasonable). Much of the knowledge that people have of their environment they get through perception; our ability to understand, predict, and explain their behaviour will depend, therefore, on our ability to attribute perceptual knowledge reliably to them.

to argue that our concepts of the senses are interest relative, not concepts of kinds of things. I am not, though, sceptical of mental kinds in general.

⁹³ It might also explain why we are interested in the fact that we ourselves have different senses: it provides us with knowledge of what we have to do in order to acquire some piece of perceptual information in the way described by Roessler in the passage I quoted in chapter 6 (see Roessler 1999). As we shall see, it is third-person ascriptions that are fundamental to explaining the distinction between the senses. This is so partly because one could have a purely practical grasp of what to do in order to acquire a certain kind of perceptual information and hence lack reflective understanding of differences between the senses; and because it is only via the third-personal interest that we can explain why we distinguish five senses in the way we do.

Robert Gordon (1995) describes a situation which involves just this kind of explanation:

You and a friend are hiking up a mountain trail, talking. Suddenly, in mid-sentence, your friend stops in his tracks, blurts out, 'Go back!' then turns and walks quietly and quickly back down the trail. You are puzzled...[t]hen you spot it: above you... a large bear, and it's a grizzly! So that's why he suddenly turned back... (p.102).

What your friend did makes sense to you once you know that he has seen, and presumably recognised, a grizzly bear in front of him. You make sense of his behaviour by attributing to him perceptually based knowledge of the bear. Although this example involves a somewhat unusual situation, we shouldn't conclude that the kind of explanation it involves is anything other than commonplace; it is, in fact, ubiquitous. In our everyday interaction with other people we are all the time judging what they know or don't know by judging what they have or have not, can or cannot, perceive. It has even been suggested that our understanding of one another as having minds depends upon or consists in our ability to know or appreciate that someone else has such a perceptual perspective on the world.⁹⁴

Our interest in what perceptual knowledge others have is not restricted to an interest in explaining their behaviour. Other people's testimony is a source of knowledge. We can come to know things through people telling us what they have perceived, but we don't treat all sources of such testimony as on a par: some people are, in certain contexts and about certain subjects, able to perceive things more reliably than others; and some things can be perceived more or less reliably depending on how they are perceived. We may take this into account when deciding what credence to give a piece of testimony. During the course of his fieldwork among the Gnao of Papua New Guinea, Gilbert Lewis found that although

it is easy to walk through the forest, there are no perspectives, no open views... The light is dimmed and greenish... The Gnao people are alert to smell... in some cases they use scent to decide the identification of trees or shrubs... The canopy and confusion of trees alters sounds and calls, limiting and muffling them, but as though enclosed in a leafy hall, the sharp screech of squawks from a nearby bird sound echoes in one's ears. I found the localisation of forest sounds difficult, ... although the native people were accurate in pointing to the direction and finding them. They excel in identifying bird calls (Lewis 1975, p.46).

If someone tells us that there is a bird of a particular species in some part of the forest – a bird that we have not ourselves perceived – we are likely to give credence to her report if she is indigenous to the forest and expert in recognising bird calls. But, given her general reliance on the distinctive call of birds to identify them together with the poor light in the forest, if were we to discover that she didn't hear the bird, but only glimpsed it through the trees, then we might doubt the reliability of the report.

Knowing what others know also plays a role in communication. In order to talk about something in the environment, it is often necessary that both the speaker and hearer know what it is that is being referred to, and know that the other knows. Thus, much communication in a shared environment seems to depend on shared knowledge of things in that environment, knowledge that is usually perceptually based.

It is often, therefore, important or useful for us to know what it is that someone else knows, or to know whether they know something. Such knowledge is useful to us not least because we use it to explain and understand other people's behaviour; because it plays a role in assessing the testimony of others; and because it is important for communication. To the extent these things are important and

⁹⁴ A suggestion that has certain affinities with the role joint attention is claimed to have in explaining an infant's conception of other minds. See, for example, the papers in Dunham, Bruner and Moore (1995).

useful to us, it is important and useful for us to know what and whether others know, where the knowledge in question is perceptually based knowledge.

If we are interested in the kinds of knowledge that people have, and if the kinds of knowledge that people have is determined by the way that they perceive things then we will be interested in which way they perceive things and we will have a reason to distinguish the different ways that they perceive things. Why is there a connection here? Because often we can judge what someone knows on the basis of judging what they can perceive.

I have argued that it is important or useful for us to know what it is that someone else knows (or does not know) when this knowledge is often determined by what that person perceives and the way that they perceive it. If we are to make use of this knowledge we must be able, in some way or other, to attribute it to others. On what basis do we make these attributions of perceptual knowledge? There are likely to be different answers to that question depending on what kind of perceptual knowledge we are attributing and for what purpose. Much of our understanding of other people is not achieved through an explicit expression by them of what they know. Sometimes we may attribute knowledge to someone on the basis of their subsequent behaviour, but not all our attributions of perceptual knowledge can be like that: often our understanding of someone's behaviour will depend on a prior assessment of what they know or perceived. This is especially true when we make judgements about whether someone did something intentionally; when we judge, for example, that he didn't intentionally hit her: he didn't see her there.

Given the connection between what someone knows about their environment and what they perceive in which way, we could attribute knowledge to others on the basis of judging what things they perceive and in which way; and surely this is often just what we do. We are often able to tell, simply by observing what they do, what someone can perceive and in which way they perceive those things. That is, we can tell what other people can see, what they can feel, hear, taste, and smell. Exactly what grounds our ability to tell these things is an interesting question. We can often tell what someone can see simply by following the direction of their gaze, for example, and this is an ability which develops at an early age in children. We can often tell what someone is able to feel by assessing what things

they are touching and the way they are touching those things: did they, for example, merely brush against the thing or did they run their hands over it. I'll say something more about how we make these attributions below. What's important to stress here is that we can make attributions of perceptual knowledge – that it is possible to do so – on the basis of what we can observe people to do.

It might be objected that, although what knowledge someone has is partly determined by the way that they perceive, that doesn't necessarily give us a reason to be interested in the different ways that they perceive things because there could be other ways of finding out about what they know. We would have a reason for as long as we sometimes attribute knowledge in this way, and surely it is the case that we both can and do determine what others know by attending to the way they perceive things.

This provides an explanation, I suggest, of the reason we distinguish different senses. Given our interest in the perceptual knowledge people acquire and the connection between the way that people perceive things and what knowledge they acquire, we have an interest not just in whether someone has perceived something, but in which way they perceived it. We have a reason, therefore, to mark a distinction between the different ways that people perceive things.

I have argued that we can attribute perceptual knowledge to other people on the basis of judging what they perceive and in which way. Because of the connections between what they know and the way that they perceive, there are various kinds of understanding of other people we can achieve through noticing which way they perceive things. That, gives us a reason to mark a distinction between the senses in as much as it gives us a reason, when attributing perceptual knowledge to others, for being sensitive to which way they perceive things but, it might be said, that doesn't explain why we should have concepts of different ways of perceiving, concepts of different senses. We could be sensitive in this way if we simply have some primitive ability to tell what other people can perceive;⁹⁵ there's no reason to think that such an ability must consist in a reflective understanding of

⁹⁵ In fact, there is some evidence that this is a very primitive ability acquired by very young children. These children are not reasoning, in any normal sense of that word, about what people can perceive. For discussion, see, O'Neill, Astington, and Flavell 1992, and Gopnik 1999.

the connection between the different ways that people perceive, and the knowledge that they acquire. Such sensitivity might manifest itself simply in an ability to correctly attribute perceptual knowledge to others with no understanding or thought about how we do that.

I want to explain what the point of our having *concepts* of the different senses is; an explanation of why we should be sensitive in this way amounts to less than that. If I want to explain the point of our having concepts of different ways of perceiving then I need to explain the point of having a fully explicit understanding of the connection between the different ways that people perceive and what they know, rather than the kind of primitive sensitivity described.

Being merely sensitive to the different ways that people perceive could explain our attributions of perceptual knowledge during our, so to speak, real-time or ongoing interactions with one another, but it cannot explain how we attribute perceptual knowledge to people at other times. Often we make attributions of perceptual knowledge to people at the time the person to whom we are making the attribution acquires that knowledge; that is, we make attributions on the basis of what we can tell about which way someone is currently perceiving something. But there are many other occasions when we could attribute some piece of perceptual knowledge to someone but, for whatever reason, have no interest in doing so at the time they acquire that knowledge. Nevertheless, at some later time we might acquire an interest in attributing knowledge to them, and we may do so at a time when we are no longer in a position to do so on the basis of observing which particular way they perceive something. At that later time we will still be able to make an attribution for as long as we remember which ways they perceived things earlier.

You might, for example, be engaged in talking with a friend sitting in a bar. Your attention and thought may be entirely taken up with the ongoing discussion. What aspects of your shared environment your friend is aware of, beyond the immediate concerns of the conversation, may be of little interest to you. The question of whether she is aware of anything else may simply not occur to you: it is not a question that is relevant to your ongoing interaction. That is not to say that it couldn't be. If your companion were to suddenly turn away from you and hide her face you might be puzzled and wonder why. Suppose that you then see someone

who you know she wants to avoid and who must have just came in the door that is behind you. You can tell that she can see the door over your shoulder and so you explain or understand her behaviour as the consequence of her seeing and recognising someone who she wants to avoid. The question of what and whether she is aware of things may in this (or some other) way become relevant; something may happen which makes it so – in this case her puzzling behaviour. But equally, the question may become relevant much later, perhaps in an entirely different context. Alone in your room you might wonder whether your friend knows that Paul has just returned from his trip abroad. What you expect her to do depends on whether or not she knows that. You know that he has because earlier you saw him put his head in at the door of the seminar room in which you were both sitting; but since your friend was sitting with her back to the door you conclude that she didn't see him and so doesn't know. Or, to pick a different example, the location of a lecture has been changed and you wonder whether or not your friend knows where it's going to be. This question hasn't occurred to you before, but now that you come to think of it, you remember seeing her looking at the poster advertising the change, and so conclude that she will most likely know of it.

There are many cases like this where we draw on our knowledge of what people have seen or not seen in now attributing knowledge – of various kinds – to them. There are other occasions when the fact that somebody perceived something in one way rather than another is important, just as it is in our real time interactions with one another. We may, for example, remember that someone saw something but didn't feel it, and so later draw conclusions about what they are likely to know or not to know as a consequence. "Why," we might ask, "does he think that he can move it without help?" – "Because, although he's seen it he has not felt how heavy it is." Just as in our real-time interactions with people, so with other explanations which appeal to what someone knows; keeping track of the different ways in which people perceived things is a way of determining what they know. I leave the reader to think up other examples, of which there are many.

Keeping track or remembering the ways that people perceive things goes beyond the kind of sensitivity that could explain our real-time attributions of perceptual knowledge. In order to attribute knowledge about what someone has perceived, we need an explicit understanding of the relation between what they did,

and what knowledge they are likely to have acquired. Only if we have such an explicit understanding could we engage in the patterns of reasoning that I have described. Such patterns of reasoning require that we conceptualise – distinguish in thought – different ways of perceiving. We can explain the point of our concepts of the senses, then, as arising out of an interest in what perceptual knowledge other people have.

If something is useful or of interest to us, then it seems natural that any community would want to talk about it. So if I am right that making distinctions between different ways of perceiving – between different senses – and keeping track of or remembering these ways of perceiving is useful or of interest to us, then it's natural that we should want to talk about ways of perceiving. This, I suggest, explains why concepts of different senses should have entered into our everyday folk psychological conceptual repertoire and vocabulary.

In general, given the connection between the particular way that someone perceives something and the kind of knowledge that they are likely to have acquired, it can be useful to track or remember the way that they perceived it. If you remember the way that someone perceived something, then you can, at a later time, work out what they are likely to have perceived, and so what perceptual knowledge they are likely to have or lack. You can do this even when, at the time you observed them to be perceiving in that particular way, you were not interested in what knowledge they acquired and at that time made no judgement about what they know. We can, therefore, think of keeping track or remembering the particular ways that someone perceived something as an especially useful way of keeping track of or remembering what they are likely to know. For as long as you remember the way someone perceived something you can, at a later date, and at a time when knowing it is important or useful, work out what knowledge they are likely to have acquired or to have failed to acquire.

We have an interest in the different ways that people perceive things and we have reason to keep track of or remember the particular ways that they perceive things, and hence to think about and have concepts of different ways of perceiving. This, I suggest, is sufficient to explain the point of our having concepts of different senses: it explains the utility of making distinctions between different senses, and of

having concepts of those distinctions. It's now time to link this to the distinctions that we make between the *five* senses.

I want to explain why we distinguish five different senses. In previous chapters I argued that there are not five different sensory capacities or mechanisms which could explain that. Now, it is important to note that the explanation that I have sketched of the distinction that we make between different ways of perceiving – different senses – doesn't depend on the existence of distinct sensory capacities corresponding to the different ways of perceiving that we distinguish; in particular, it doesn't depend on the existence of five capacities corresponding to the five senses.

We can observe people to do different things with respect to the things around them, and so we can judge which things around them they can perceive and in which way they can perceive them. So, for example, we can tell in which direction a person's eyes are directed, what things they are touching with their hands, whether they have put something in their mouth, whether a sound is loud enough for them to perceive, whether there is enough light, and so on. On the basis of this we can attribute perceptual knowledge to them. Making attributions on this basis works – is reliable – because there is a connection between what someone does – and can be observed to do – with respect to the things around them and the particular ways that they perceive those things, and so what knowledge they acquire. This is true whether or not there actually are distinct perceptual capacities corresponding to the different ways that we judge people to perceive things. For as long as differences in the way that people interact with the things around them leads to differences in the kind of perceptual knowledge they acquire, it doesn't matter to the truth of our attributions what the correct explanation of that is. In particular, it doesn't matter what the perceptual mechanisms or capacities in virtue of which that is true are. The different ways that we judge people to perceive things do not need to correspond to, or pick out, any underlying capacities or mechanisms. We are interested in what knowledge people are likely to have, and as long as we track what they do we can succeed in attributing knowledge to them: all that is required is that there exists a reliable connection between the different ways that people can be observed to interact with the things around them, and what perceptual knowledge they are likely to acquire. The existence of such a reliable connection may well

depend on the existence of various different kinds of sensory capacity whose exercise is manifested in the different ways people interact with the things around them, but we don't need to track those capacities in order to make reliable attributions of knowledge. All we need to track is the ways that people interact with things, and to make attributions of the basis of observing the different ways that people interact with the things around them.

If the point of the concepts can be explained in terms of an interest we have in tracking the way that people interact with the things around them for the purposes of attributing perceptual knowledge to them, then it doesn't matter whether or not our concepts of different senses correspond to, or pick out, any independently existing perceptual capacities. In fact, there's no reason to think that the particular ways of perceiving that we track, and have concepts of, are the only ways of perceiving that it could be useful to track.

When we make attributions of perceptual knowledge to other people we can do so on the basis of making more or less careful observations in an attempt to determine with respect to what they have satisfied the conditions necessary for perception. Our grounds for making attributions can be more or less fine-grained, more or less specific, depending on what we want to know about what they know. We might, for example, want to know what someone knows about the things that they perceive, or we might be interested simply in whether someone has perceived some object or other particular thing. Sometimes we might want to know whether someone has some quite specific piece of perceptual knowledge, whether they have perceived some particular state of affairs or fact: "Has she seen that the floor is wet?", "Did he hear what she said?", and so on. We want, in such cases, to know of someone whether they perceived that such-and-such is the case. At other times we might want to know what someone perceives of their immediate environment more generally; what, of the things around them that they could have perceived, have they in fact perceived? This kind of knowledge can be useful in explaining someone's behaviour: "Why did he storm off like that?" – "Because he overheard us talking about him," or "Why did he suddenly walk away?" – "Because he saw a bear." At other times we might be interested, as in the example of bird recognition I described earlier, in how reliable someone's claims to knowledge are, in whether they really were in a good position to recognise something.

In attempting to make these different kinds of knowledge attributions we need to pay attention to different aspects of what someone does with respect to the things around them, and the way those things around them are, in order to tell which conditions necessary for perception have been satisfied. We might simply pay attention to what things they are touching with some part of their body, what they put in their mouth, which direction their eyes are looking, whether some sound is loud enough for them to hear, and so on. But sometimes our attribution of perceptual knowledge to someone might be based on a judgement that they are doing some quite specific thing. There may, for example, be some contexts in which it is important for us to know whether someone has perceived the colour or shape of something. In that case we have a reason to base our judgements not just on whether they looked at the object, but whether they have done so in a way that would make it likely that they perceived the particular property in question, whether they have done some specific thing which would make it very likely that they perceived the object's colour or shape. Of course it's usually (although not always) true that if one perceives an object by looking at it then one perceives its colour and shape, and that if one perceives the colour of some object then one also perceives its shape, and *vice versa*. But if it is important for us to know whether someone really can perceive the colour of something we might attempt to determine whether they are looking at it in good, neutrally coloured, light; whether it is sufficiently near to them and oriented appropriately; that it's not a specular surface whose colour is liable to be obscured by reflections; and so on. If, on the other hand, it was important for us to know whether they can perceive the object's shape, we might attempt to determine whether some other conditions relevant to their perceiving that kind of property are met.

That there are such differences in what conditions have to be satisfied in order to perceive different properties of objects is more obvious with respect to touch. Which features of something one can perceive by touching it depends on how one touches it. It depends for example, on whether one actively explores some object with one's hands or whether the object is simply pressed up against one. If one actively explores an object with one's hands what one perceives, what one can tell about it, depends on whether one presses, prods, rubs, or hefts it. What one perceives by touching depends, too, on the part of one's body one touches the

object with. Different parts of the body are differently sensitive to different properties of things.⁹⁶ If we want to know whether someone has perceived some particular property of something that they are touching with a part of their body we may pay attention to the quite specific kind of way that they interact with it.

If the things to which we pay attention – the conditions necessary for perception we judge to be satisfied – in deciding what perceptual knowledge someone has can be more or less specific, then, were we to classify them, we could do so in a variety of different ways. I argued above that we have a reason to track or remember the ways that people perceive things; if what I have just argued is right, then what we have a reason to track or remember is the kinds of ways people interact with the things around them. Our concepts of the senses are just concepts of these ways that people interact with the things around them. For as long as we have an interest in keeping track of or remembering the ways that people interact with the things around them then we will want to classify them in some consistent way; it's only then, for example, that we will be able to talk to one another about them. We classify the ways that people interact with things in five ways – we have concepts of five ways of perceiving, five senses. What explains why we classify them in this way? Why do we distinguish five senses? Given the point that making distinctions between the senses has for us, we needn't suppose that an explanation of why we distinguish five senses must appeal to some real distinction between different sensory capacities. We can explain why we classify them as we do as being a matter of convention. We needn't suppose that there is any correct classification of ways of perceiving in order to explain why we have the concepts of the senses that we do.

We have concepts of five senses, and although I think we can go some way to explaining why we have just five, and why we have just the five that we do, it is ultimately a contingent and arbitrary matter that we make the distinctions in the way we do. We could have distinguished different senses, or distinguished the senses

⁹⁶ For a further discussion of perception by touching see Gibson (1966, Ch.VII) and see Lederman and Klatsky (1987 and 1990) for a discussion of the hand movements involved in tactual object recognition.

differently, and we could have done so even had nothing else about our senses – their mechanisms or the kind of experience that we have – been different.

3. What does it mean to say that it is a convention that we classify the senses as we do? There may be many habits or regularities of behaviour that are conventional, some are obviously so: using a certain spoon to eat the soup course, driving on the left hand side of the road, or meaning something determinate by a particular word. What makes these kinds of regularity in behaviour conventional regularities? It is not that those who conform to them entered into an agreement to do so, since there are many conventions whose participants have never entered into an agreement. I suggest that it means, rather, that the distinctions we make are arbitrary and that were it not for some historical accident, we could have made different distinctions to roughly the same purposes (c.f. Burge 1975, p.249). What is it for a convention to be arbitrary? A convention is arbitrary, Burge suggests, if

[a]s a matter of fact – whatever the participants may believe – it is within the power of the participants to have learned an incompatible regularity that would have served substantially the same social functions without demanding significantly greater effort on the part of the participants (p.254).

This captures the sense in which – as I want to claim – the way we classify the five senses is arbitrary. There are, on this view, two aspects to the arbitrariness of a convention. The first is that the regularities which are conventional are not determined by biological, psychological, or sociological law, so the fact that someone learns a particular convention is a matter of historical accident. The second is that the regularities which are conventional are not uniquely the best means of fulfilling their social functions: other incompatible means would have done as well (p.254).⁹⁷

⁹⁷ I take social function to be social not in the sense of being specific to a particular society or culture, but as being a practise that essentially involves others. The social function that I have claimed the senses to have will be the same across different societies.

What reason is there to think that the classification we make of the senses is arbitrary in this way? We don't think they are arbitrary, but that we don't doesn't imply that they are not, since the fact that we currently lack a reason to believe (or even to disbelieve) that a given regularity is a convention does not preclude us from deciding later that it is and was such. In his well known account of conventions, Lewis stresses the rational basis of conventions and he succeeds in showing why it is often reasonable to participate in a convention (Lewis 1969). According to him, a regularity is conventional for a population if and only if most members of the population conform to the regularity and (at least an important part of) the explanation of why they do so, as opposed to conforming to any other equally serviceable rival, is that they each expects others to do so, and each prefers to do so if the others do. Lewis, however, "makes it seem as if, were the parties to a convention irrational in their actual motives, overly insistent on a particular means, or insufficiently intent on their recognised end, there would be no convention" (Burge 1975, p.252). Often the parties to a convention are confused about their relevant ends, and the social functions of their practice; they are brought up achieving them and do not know the origin of their means; and they may disagree about whether another means is possible or they may simply fail to consider the question. It would not be an objection to my claim, then, that the way we classify the senses is conventional that we may not know of any alternatives, and that we may not believe that our classification is arbitrary.

Distinguishing and keeping track of ways of perceiving has a function, that of allowing us to attribute perceptual knowledge to others of the kind required for various kinds of understanding and communication. But this function has a social value, too, in as much as it allows us to talk about and communicate what we know about other people. So, just as there is a value for each of us in keeping track of the ways in which other people perceive, so there is a value for each of us if others keep track of the same kinds of ways of perceiving. It's this which makes it plausible to claim that we keep track of the ways of perceiving we do as a matter of convention.

I do not deny that there are different ways of perceiving things; I claim that the senses just are five different ways of perceiving that we have names for and concepts of. But, as I have argued in previous chapters, there is no natural or interest independent way of drawing distinctions and categorising these different

ways of perceiving into five; our categorisation of ways of perceiving into five is therefore arbitrary. We could have had names for, and have kept track of, ways of perceiving other than those we do, which means that that we could have had more or fewer senses than we actually have, and we could have had different senses to those we actually have.

Someone who denied that our classification of ways of perceiving into five senses is arbitrary might insist that the classification we actually make constitutes uniquely the best means of fulfilling the social function that distinguishing different senses serves and so is not arbitrary. If this were true, the distinctions couldn't have been different and still "served substantially the same social functions without demanding significantly greater effort on the part of the participants". The best way of answering this objection is to indicate how an alternative way of distinguishing the senses – an alternative classification – might serve the same social function as the classification we actually make. We can show that if we can suggest why we might have arrived at the classification of the senses that we have, and describe the kinds of circumstances which might have led us to make a different classification. We could divide up the ways that we judge people to perceive, and on the basis of which we can make attributions of perceptual knowledge, in almost any way we chose. Why do we do it one way – the way that we actually do – rather than another? Part of the point in claiming that the classification of the senses we make is arbitrary is that there need be no explanation of why we classify the senses as we do, it's simply an historical accident. But it is not the case that our classification is completely without constraint, and I think that we can say something about why we should have settled on the classification of ways of perceiving that we have, and the kinds of circumstance that might have led – might lead – us to make alternative classifications. It is possible to describe the kinds of circumstances which we can speculate could have led to different ways of distinguishing and classifying the senses.

Given that we are interested in classifying the senses because we are interested in what knowledge people acquire, we will only have a reason to distinguish ways of perceiving that lead to the acquisition of different kinds of knowledge. We would have no reason, for example, to distinguish what someone does with their left-hand from their right-hand, their left-eye from their right-eye.

Another important part of any explanation of why we classify the senses we do is the ease with which we can tell that someone is perceiving something one way rather than another. Although it is better to be more specific in our judgements about the ways someone perceives something, because then we will be more precise in our attributions of knowledge, and so better at understanding them, there is a cost in being too specific. It is hard to tell whether someone really is perceiving that way; it is, for example, relatively easy to tell whether or not someone is looking at something – a painting, say – harder to tell whether they are in a position to perceive its colour, harder still to tell whether they can perceive the small scratch in the bottom corner of the frame. It will be difficult, too, to keep track of and remember a large number of different ways of perceiving. It is sometimes said, for example, that there are in fact many more senses than we have names for (Rivlin and Gravelle (1984) count seventeen senses; what they count as a sense is something physiological) but there would clearly be little benefit to be gained from having concepts of many more than we already have. We gain through generality and simplicity what we lose in specificity, but there is no non-arbitrary way of weighing gains against losses.

It might be objected that the ease with which we are able to tell that someone is perceiving one way rather than another is not arbitrary. In particular, given that our judgements will be based on what we can observe people to do, that certain parts of the body are involved in perceiving might be thought to determine which ways of perceiving we track and so how we classify the senses. That we judge someone to have perceived something using a particular part of the body clearly is important to the distinction as we make it. We judge someone to have seen something only if they perceive it in a way that involves using their eyes; we judge someone to have felt something only if they perceive it by touching it with some part of their body, in particular with their hands. With the other senses things are a little more complicated. Although we don't judge someone to have tasted something unless they perceive it by putting it in their mouth, we judge someone to have tasted something only if they perceive its flavour by putting it in their mouth (we can make sense of something being perceived to have no flavour when put in the mouth). With hearing and smell the role of particular parts of the body seems less important. Of course we think that hearing involves the ears, but we don't

judge someone to have heard something only if they perceive it with their ears: we cannot tell whether they did perceive it with their ears independently of judging them to hear it, and we would judge that someone without ears (without those external parts of the body) could still hear things; the same applies to smelling. For the same sorts of reason we don't judge someone to have smelled something only if they perceive it using their nose. Parts of the body are important to the distinction we actually make, I suggest, because we can easily tell that they are used, and because perceiving things in a way that involves different parts of the body leads to the acquisition of different kinds of knowledge. But that doesn't mean that the distinctions we make between different senses and our classification of them is determined by the structure of our bodies,⁹⁸ since the extent to which that determines the way we classify ways of perceiving can be offset or modified by other kinds of consideration, some of which vary in an arbitrary way.

It seems, too, that the more likely the perception of some feature is to occur in the absence of the perception of any others the more likely we are to distinguish the conditions for the perception of it. Why bother, for example, to distinguish what someone must do to perceive colour from what they must do in general to perceive something with their eyes, given that the two reliably go together? And given that many cases of perceiving something by contact or touch are not sufficient for the perception of the flavour of that thing, we should expect the conditions required for the perception of flavour, as opposed to merely perception by touch to be classified as distinct. That wouldn't be true, though, if the flavours things had were a matter of indifference to us. We should expect our differing interest in different kinds of things we can perceive to play some role in explaining why we distinguish and classify the senses that we do. If it is a matter of indifference to us whether someone can perceive a certain kind of thing, then we will be indifferent about judging whether or not they can perceive it; conversely, if it is very important for us to know whether someone has perceived some kind of thing then it will be important for us to be able to judge whether they have. The more interest a particular thing has for us, the more likely we are to want to know that conditions

⁹⁸ By parts of the body I don't mean sense organs. Remember that what we regard as a sense organ is to a certain extent arbitrary, arbitrary for the same reasons the senses are arbitrary.

necessary for its perception have been satisfied, and so the more likely we are to distinguish those conditions from others, and classify them as a distinct way of perceiving.

Whether it is important for us to know whether or not someone can perceive something is the kind of thing that will vary from culture to culture, and from environment to environment. One point that emerges from discussions of the relative value of the senses and their objects across cultures (see below) is that senses which are important for practical purposes may not be important culturally or symbolically, and vice versa. The cultural, symbolic, or mythical importance of senses and their objects can, therefore, be another source of variation which could have led to an alternative classification.

We can, then, give some explanation of what might have led us to classify five senses in the way that we have. Equally, we can imagine different, alternative, classifications we could have made. We might, for example, have distinguished perceiving something by contact from perceiving something in any other way; we might, that is, have distinguished a contact sense from a distance sense. There would be nothing intrinsically wrong in making this distinction. We might not have distinguished smell from taste, but simply grouped what are sometimes called the chemical senses together. We might have distinguished different kinds of perception by touch: hand touch and body touch, perhaps, or active and passive touch. In the right circumstances, we could have classified the senses in these ways: that we don't is, therefore, historical accident.

4. Perhaps the best evidence that a practice which serves some social function is conventional for a population is the existence of a different population who serve the same function by means of a different practice. The best evidence that the way we distinguish the senses is conventional would be the existence of a population that distinguished different senses, or distinguished senses in a different way. Are there any such populations? There is no definitive answer to that question: partly because no one has been particularly concerned to look, and partly because what evidence there is is equivocal.

Anthropologists, it seems, have not been especially interested in whether other cultures distinguish senses in any way differently to us, perhaps this is because

they take it for granted that the senses are not conventional and so don't recognise the possibility of alternatives; and perhaps, too, because that another population does distinguish the senses differently is, unless one is alive to the possibility, a thing easily missed or overlooked.⁹⁹ Even a recent collection of papers which, according to its editor, contributes to a new field of study – 'the anthropology of the senses' – "concerned with how the patterning of sense experience varies from one culture to the next in accordance with the meaning and emphasis attached to each of the modalities in perception" (Howes 1991, p.3) rarely addresses the question of differences in the number or divisions between senses across different cultures, and concentrates rather on differences in value placed on each of the five senses. In a paper that sets out the methodology of this new field of study, David Howes says that

It is this idea...with which the anthropology of the senses is most centrally concerned. In approaching other cultures, what we want to find out is: What *is* the relative importance and meaning of the different senses to the members of that culture? (pp.168-9).

Anthropologists of the senses (in this book at least) are thus interested in "the differing combination of *the five senses*" (my emphasis) in individuals from different cultures (p.168). If this is the question that they ask, then one can understand how differences in the way the senses are distinguished can go unremarked. It appears that they are simply assuming that people from other cultures distinguish five senses, as we do, and that any differences will be in the different relative value they give to each of them.

In fact it is not clear that, even when describing the relative value that different cultures place on different senses, these anthropologists really succeed in describing different valuations of the senses, rather than differences in the value given to certain kinds of things or features of things. One writer, for example,

⁹⁹ A search of Anthropological Literature on Disc (Anthropological 1998) using the keywords 'perception', 'senses', and 'five senses' returns no results relevant to this question. With the

describes the importance of, and high value placed on, scent and by implication the sense of smell, in Moroccan culture (Griffin 1991). But given that they are distinct, it doesn't necessarily follow from the fact that scent is culturally valued that the sense of smell is; that depends on what scent is valued for and why. Griffin says, for example, that "[s]trong-smelling substances are burned when one wants to keep the *jinn* or spirits away" (p.211). It is difficult to know whether we should conclude from this that the sense of smell is thought important, or whether the smell itself is important because it is thought to have special causal powers independent of the way it smells. It may be, of course, that the spirits stay away because the smell smells bad (and given the way we often tend to think of the smell of something as just the way it smells this may seem just obvious) but that needn't be the case. My point is simply that in describing these differences one must be sensitive to such distinctions; the apparent lack of sensitivity in many of these anthropological descriptions means that we should place little evidential weight on them as descriptions of actual alternatives when arguing that the way that we distinguish five senses is conventional.

Nevertheless, Howes does suggest that "[o]ther cultures do not necessarily divide the sensorium as we do. The Hausa recognise two senses; 'the Javanese have five senses...which do not coincide exactly with our five'...there may be any number of 'senses', including what we would classify as extra-sensory perception – the 'sixth-sense'" (Howes and Classen 1991, p.258).¹⁰⁰

exception of Howes (1991), and the items listed in the bibliography of that book, I have been unable to find any discussion of cultural variation of the senses.

¹⁰⁰ The sentence Howes quotes here is from Dundes (1980, p.92). Dundes's paper discusses the priority of vision in American culture, especially as reflected in metaphor. He sees as one consequence of this that, in looking at other cultures, "we run the risk of imposing our own rank-ordering of the senses upon data that may not be perceived in the same way by the people whose cultures are being described" (p.91), and he goes on to say that "Clifford Geertz reports, for example, that the Javanese have five senses..." (p.92), but gives no reference, and ends by commenting that "Cross-cultural comparisons of sense categories may...reveal critical differences in the specific senses". The implication is clearly that such cross-cultural comparisons have, at the time of writing, yet to be made. Other than those cited in the main text, I have been unable to find any more recent examples of cross-cultural comparisons of the senses.

Hausa is a language spoken by people who traditionally inhabited an area in northern central Nigeria called the Hausa Bakwai. Hausa differs from, for example, English, in that it

has only one verb for all the non-visual senses. Thus, the functions of hearing, tasting, smelling, and touching are all designated by the word *ji*. The verb *ji* also means to 'feel' things in an intuitive or emotional sense... The Hausa word *gani* means 'to see'... this word only means 'to see'. It is never used in the sense of understanding what another person means (Ritchie 1991, p.194).

It is, Ritchie suggests, an implication of this "that the Hausa recognise only two senses compared with our five – that is, a multimodal sense (the *ji*-complex) and a monomodal one (sight)... Of course, it is normally possible to tell from the context as well as the rest of a sentence in what sense *ji* is being used" (p.194). He offers no real explanation of why there should be this difference, except to suggest that vision is less important to the Hausa than the other senses and that, in the Hausa language, 'understanding' is associated with hearing, tasting, smelling and feeling, and distinguished from sight.

It would be a mistake, I think, to place much weight on this example. It is just not clear what is going on; in particular, it is not clear from Ritchie's description whether the Hausa distinguish any other senses in addition to these two (although the implication is clearly that they don't). It may be that further distinctions are made within the '*ji*-complex', but that the Hausa think there is something importantly similar or common to just those senses that comprise it – they may, for example, be thought to be a source of a certain kind of understanding or knowledge. The Hausa would then have a reason for thinking that sight and *ji* are fundamentally and importantly different in kind, but not that they are the only kinds of senses there are, since there could still be further distinctions within these two fundamentally different kinds.

What conclusion are we entitled to draw from this rather sparse anthropological evidence? We might suppose that the Hausa really do only distinguish two senses. That would be the strongest possible support for my claim

that the way we distinguish the senses is arbitrary and that we could have distinguished them differently, since we could have distinguished them as the Hausa do. But even if one is sceptical that Hausa represents an actual alternative way of distinguishing the senses, the coherent possibility that it does so again supports my suggestion, since what is possible for the Hausa is possible for us: we could have distinguished them as the Hausa possibly do.

It might be objected that to draw this conclusion is question begging: it would not be legitimate to describe the Hausa as even potentially representing an alternative way of distinguishing the senses. They have just the senses that we have, but they don't have names for them. The five senses, this objection continues, are really distinct independently of how we categorise them, they are not something that *we* have to distinguish, so what I describe as distinguishing different senses should really be described simply as lacking concepts for (some of) the senses that they, and we all, have.

I think that this objection is well made against this particular example, but I don't think that it will succeed as general argument against using evidence of cross-cultural differences to support the claim that the distinctions that we make between the five senses is arbitrary. To be convincing this kind of objection must explain why there are differences in the concepts of the senses across cultures despite there being no differences in the senses, and there will come a point, I suspect, when such explanations will be less plausible than the claim that the distinctions between the senses are themselves arbitrary.

It might be objected that the way that other cultures distinguish different senses can never be evidence for the claim that the way we distinguish senses is arbitrary. It can never be such because there is no reason to think that the people from other cultures that we come across really do have different senses to us, rather than that simply conceptualise the – same – senses differently to us. Why should we think that they have the same senses as us? Someone who makes this objection might appeal to the fact that they have the same sensory physiology as us, and that they have the same practical grasp of the distinctions between the senses that we have (manifested perhaps in the kinds of practical knowledge they have of what to do in order to perceive something). I have argued that these similarities do not imply that we have the same senses, but the objector claims that I cannot use

anthropological evidence as evidence for that argument without begging the question. I want to give an explanation of our common sense, folk-psychological, concepts of the senses. If there exist cultures with different folk-psychological concepts of the senses then that suggests that either one or both cultures is mistaken about what senses they have, or that the concepts are not concepts of independently existing senses. Given that I take it for granted that our judgements involving concepts of the senses are not globally false, I claim the latter. Given that assumption, to make such a claim is not question begging. Someone who makes this objection must claim that one or other or both of these different ways of conceptualising the senses is wrong. It is difficult to see what justification they could have for such a claim.

5. According to my account, then, the five senses are just a conventional way of classifying the different ways in which people can perceive things. When we judge someone to perceive something with a particular sense we judge them to perceive in one way rather than another. That these different ways get grouped together into five senses is a matter of convention and, to a certain extent, arbitrary: we could have classified them differently. Our classification of ways of perceiving is not determined by any interest independent distinction amongst ways of perceiving.

Someone sees something, on my account, just in case they perceive it in that way we classify as seeing; someone sees rather than feels the shape of something that they both see and feel just in case they perceive the shape in the way we classify as seeing, rather than in the way we classify as touching. In chapter 2 I said that an account of the senses must explain why, and on what basis, we apply concepts of the senses to the various activities that we and others engage in when perceiving, and that an account must explain why we apply concepts of the senses to various experiences. The account I have described here can provide such an explanation. In this chapter I have been arguing that we apply concepts of different senses to people by observing what they do, and that it is this which explains why we distinguish the senses as we do. We can explain why we apply concepts to kinds of experience by supposing that we classify an experience of something as, say, a visual experience if we see the something in virtue of having that experience.

My account of the way that we distinguish the five senses has clear affinities with the kind of account which appeals to a simple theory of perception. It differs from that account in the way that it explains why we distinguish five senses as we do, and in what it claims is common to all the perceptions of a particular sense. The attempt to distinguish the senses by appealing to a simple theory of perception failed, I argued, because such a theory would not distinguish five senses: there are not five kinds of sense for such a theory to be a theory of. Merely appealing to a simple theory of perception cannot, therefore, explain why we distinguish five senses. My account explains why we distinguish five senses without appealing to the independent existence of five senses. It does so by appealing to conventions; we can appeal to conventions to explain the way we distinguish the senses because of the social role the concepts of the senses have in attributing knowledge to other people and in explaining their behaviour. Notice that I do not deny that we have different sensory capacities: I don't deny that to perceive something by seeing and by touching involves the exercise of different capacities; it is just that we cannot explain the way we distinguish *five* senses, and have concepts of *five* senses, merely by appealing to that fact. My account is able to offer an explanation, albeit a rather deflationary one, of why we have five senses rather than more or fewer. The account claims that, although we distinguish five senses there need be nothing in common to all the perceptions of a particular sense in virtue of which they are the perceptions of that sense. It doesn't need something in common to all such perceptions in order to explain why we distinguish *five* senses. The perceptions of a particular sense do not form a kind, and so our concepts of the five senses do not pick out kinds of thing. A consequence of this is that there is no straightforward answer to the question of whether other creatures have the same senses as we do.

In chapter 5 I described an account of the way we distinguish the five senses which I called the Brute Experiential View; I said that the only reason to think the Brute View is the correct account of the distinction is that there is no alternative explanation. In this chapter I have developed an alternative explanation, but what reason is there to prefer my explanation of the distinction over the Brute View? We might reject the Brute View because we don't think that experiences have features of the kind to which it appeals, but even if we don't reject it for that reason, to the extent my account offers a better explanation of the senses than the Brute View we

have a reason to accept it in preference to the Brute View. Does my account offer a better explanation of the distinction and, if so, why?

I suggested that there could and may be cultural variation in the way that people distinguish different senses, and that would provide a reason for thinking that my account is correct. It would equally give us a reason to reject the Brute View. The only way that the Brute View could explain such cultural variation would be by appealing to differences in the features that perceptual experiences have in different cultures. It can, however, provide no explanation of why there should be such differences, and why they should be culturally determined. If one is convinced, therefore, that such variation is possible, one will have reason to reject the Brute View.

Part of the reason the Brute View cannot explain cultural variation in the senses is that the explanation that it offers of the distinction is, in a way, shallow. It is just a brute fact, not susceptible to further explanation, that we enjoy five different kinds of experience, and so have five senses. That these features happen to be connected with experiences produced by various parts of the body, and that they are connected with experiences of certain kinds of thing is, again, something that cannot be explained. In a way, then, the Brute View can offer no real explanation of why we have five senses because it offers no real explanation of why our experiences should have five kinds of feature corresponding to each of the senses. In fact, I suggested, the only reason it can put forward that our experiences do have the five kinds of feature in question is that we distinguish five senses. My account, although deflationary, offers a deeper explanation of the distinction. It can provide an explanation of why we should have distinguished the senses in the way that we do, and an explanation of why we have five senses, and of the kinds of reason that might have led us to distinguish the senses differently. My account can also provide a better explanation of the role of the body plays in distinguishing the different senses, and of why the distinction that we make is so closely tied to the kinds of things that we perceive. On the Brute View, that certain parts of the body are connected with the use of different senses is a contingent matter: it just so happens that those parts of the body produce the experiences that we call visual, tactual, and so on. To many people, the connection between the sense with which we perceive and the part of the body involved does not seem contingent, and on my

account we can see why not. If one finds the kind of shallow explanation offered by the Brute View unsatisfactory, then one should favour my account instead.

These reasons do not constitute a conclusive argument against the Brute View; nonetheless, I suggest they provide reason enough to reject that view in favour of my account.

6. In this thesis I have attempted to give some account of what constitutes the five senses of common sense; to give, that is, some account of the everyday distinction that we make between five different senses. I have put forward some sceptical arguments against some accounts of the distinction and suggested that those accounts are incapable of explaining the distinction as we make it. Instead I have sketched an alternative account which, I claim, can explain the distinction as we make it, and can do so without supposing that there really are five senses independent of our classification and conceptualisation of them. Given that the senses are not independent of our classification in this way, we needn't suppose that a theory of perception has to address questions about the five senses. A theory of perception should simply be a theory of our perception of particular things.

There are many other questions that remain to be addressed. Most importantly, questions about our *knowledge* of the senses. Throughout this thesis I have been considering a constitutive question about what the senses are, and have set aside questions about how we know what senses we have, and on what basis we judge ourselves to be perceiving with one sense rather than another. Given my account of the senses, an answer to that question will appeal to contingent features of our experience; that is, the grounds on which we judge ourselves to be perceiving with a particular sense will not be what constitutes the difference between the senses. There is clearly more to be said about that question, and in general about how we come to know about our own experience. But they are questions for another occasion.

Bibliography

- Ackerman, Diane. 1990. *A Natural History of the Senses*. London: Chapman's.
- Aglioti, S., M. A. Goodale, and J. F. X. DeSouza. 1995. Size-contrast illusions deceive the eye but not the hand. *Current Biology* 5: 679-85.
- Anscombe. 1971. Causality and Determinism. In *Causation and Conditionals*, edited by E. Sosa. Oxford: Oxford University Press.
- Anthropological. 1998. *Anthropological Literature on Disc* [compact disc]. Tozzer Library, Harvard University.
- Armstrong, D.M. 1962. *Bodily Sensations*. London: Routledge and Kegan Paul.
- Armstrong, David M. 1997. *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Atherton, Margaret. 1990. *Berkeley's Revolution in Vision*. Ithaca: Cornell University Press.
- Austin, J.L. 1962. *Sense and Sensibilia*. Oxford: Oxford University Press.
- Ayer, A.J. 1973. *The Central Questions of Philosophy*. Harmondsworth, England: Pelican.
- Ayers, Michael. 1993. *Locke. Epistemology and Ontology*. London: Routledge.
- Barlow, H. B., and J. D. Mollon. 1982. *The Senses, Cambridge Texts in the Physiological Sciences; 3*. Cambridge: Cambridge University Press.
- Berkeley, George. 1732. An Essay Towards a New Theory of Vision. In *Berkeley Philosophical Works*, edited by M. R. Ayers. London: Dent.
- Berkeley, George. 1734. Three Dialogues Between Hylas and Philonous. In *Berkeley Philosophical Works*, edited by M. R. Ayers. London: Dent.
- Bodanis, David. 1986. *The Secret House: the Hidden World of Your Home*. London: Sidgwick & Jackson.
- Brandt Bolton, Martha. 1994. The Real Molyneux Question and the Basis of Locke's Answer. In *Locke's Philosophy: Content and Context*, edited by G. A. J. Rogers. Oxford: Clarendon Press.
- Bregman, Albert S. 1990. *Auditory Scene Analysis*. Cambridge, Mass.: MIT Press.
- Bruce, Vicki, and Patrick Green. 1990. *Visual Perception: Physiology, Psychology and Ecology*. London: Lawrence Erlbaum Associates.
- Burge, Tyler. 1975. On Knowledge and Convention. *Philosophical Review* 84: 249-254.

- Campbell, John. 1994. *Past, Space, and Self*. Cambridge, Mass.: MIT Press.
- Campbell, John. 1996. Molyneux's Question. In *Perception*, edited by E. Villanueva. Atascadero: Ridgeview.
- Campbell, John. 1997. Sense, Reference, and Selective Attention. *Proceedings of the Aristotelian Society, Supplementary Volume LXXI*.
- Carterette, Edward Calvin, and Morton Philbert Friedman. 1978. *Handbook of Perception, Volume VIa: Tasting and Smelling*. New York: Academic Press.
- Churchland, Paul. 1981. Eliminative Materialism and the Propositional Attitudes. *Journal of Philosophy* LXXVIII (2): 67-90.
- Churchland, Patricia. 1989. *Neurophilosophy: Towards a Unified Science of the Mind-Brain*. Cambridge, Mass.: MIT Press.
- Clark, Thompson. 1965. Seeing Surfaces and Physical Objects. In *Philosophy in America*, edited by M. Black. London: Allen and Unwin.
- Coady, C.A.J. 1974. The Senses of Martians. *The Philosophical Review* 83 (1): 107-125.
- Cowey, Alan. 1979. Cortical Maps and Visual Perception. The Grindley Memorial Lecture. *Quarterly Journal of Experimental Psychology*.
- Craig, Edward. 1990. *Knowledge and the State of Nature: an Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- Cummins, Robert. 1983. *The Nature of Psychological Explanation*. Cambridge, Mass.: MIT Press.
- Davies, Martin, and Tony Stone. 1995. *Folk Psychology: the Theory of Mind Debate*. Oxford: Blackwell.
- Degenaar, Marjolein. 1996. *Molyneux's Problem: Three Centuries of Discussion on the Perception of Forms, International Archives Of The History Of Ideas; 147*. Dordrecht, Netherlands: Kluwer Academic Publishers.
- Dretske, Fred I. 1995. *Naturalizing the Mind, The Jean Nicod lectures; 1994*. Cambridge, Mass.: MIT Press.
- Dundes, Alan. 1980. Seeing is believing. In *Interpreting Folklore*. Bloomington: Indiana University Press.
- Dunham, Philip J., Jerome S. Bruner, and Chris Moore. 1995. *Joint Attention: Its Origins and Role in Development*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Dupré, John. 1993. *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, Mass.: Harvard University Press.

- Evans, Gareth. 1982. *The Varieties of Reference*. Edited by J. McDowell. Oxford: Clarendon Press.
- Evans, Gareth. 1985a. Molyneux's Question. In *Collected Papers*. Oxford: Clarendon Press.
- Evans, Gareth. 1985b. Things Without the Mind. In *Collected Papers*. Oxford: Clarendon Press.
- Everson, Stephen. 1997. *Aristotle on Perception*. Oxford: Clarendon Press: Oxford University Press.
- Farah, Martha J. 1995. Dissociable systems for recognition: a cognitive neuropsychological approach. In *An Invitation to Cognitive Science Volume 2: Visual Cognition*, edited by S. M. Kosslyn and D. N. Osherson. Cambridge Mass.: MIT Press.
- Fodor, Jerry. 1968. *Psychological Explanation: An Introduction to the Philosophy of Psychology*. New York: Random House.
- Fodor, Jerry. 1983. *The Modularity of Mind*. Cambridge, Mass.: MIT Press.
- Fodor, Jerry. 1990. Précis of Modularity of Mind. In *A Theory of Content and Other Essays*. Cambridge, Mass.: MIT Press.
- Fodor, Jerry. 1974. Special Sciences. *Synthese* 28.
- Gibson, James Jerome. 1966. *The Senses Considered as Perceptual Systems*. Boston: Houghton Mifflin.
- Goodale, M.A., Péllisson, D., Prablanc, C. 1986. Large adjustments in visually guided reaching do not depend on vision of the hand or perception of the target displacement. *Nature* 320: 748-50.
- Gopnik, Alison. 1999. Infant's and Toddler's Understanding of Visual Perception. Paper read at Joint Attention Workshop, at HRB Joint Philosophy and Psychology Project on Consciousness and Self Consciousness.
- Gordon, A.M., Forssberg, H., Johansson, R.S. and Westling, G. 1991. Visual size cues in the programming of manipulative forces during precision grip. *Experimental Brain Research* 83: 477-82.
- Gordon, Robert M. 1995. The Simulation Theory: Objections and Misconceptions. In *Folk Psychology*, edited by M. Davies and T. Stone. Oxford: Blackwell.
- Gregory, R. L. 1966. *Eye and Brain: the Psychology of Seeing*. London: Weidenfeld & Nicolson.

- Grice, H.P. 1967. Some Remarks About the Senses. In *Studies in the Way of Words*. Cambridge, Mass.: Harvard.
- Griffin, Kit. 1991. The Ritual of Silent Wishes: Notes on the Moroccan Sensorium. In *The Varieties of Sensory Experience*, edited by D. Howes. Toronto: Toronto University Press.
- Griffiths, Paul. 1997. *What Emotions Really Are: The Problem of Psychological Categories, Science and its Conceptual Foundations*. Chicago, Ill.: University of Chicago Press.
- Hamlyn, D. W. 1975. *Aristotle's De anima: books II and III (with certain passages from Book I)*, Clarendon Aristotle Series. Oxford: Clarendon Press.
- Harman, G. 1990. The Intrinsic Quality of Experience. In *Action Theory and Philosophy of Mind*, edited by J. Tomberlin. Northridge: Ridgeview.
- Hopkins, Robert. 1998. *Picture, Image and Experience: a Philosophical Inquiry*. Cambridge: Cambridge University Press.
- Howes, David. 1991. *The Varieties of Sensory Experience: a Sourcebook in the Anthropology of the Senses*. Toronto: University of Toronto Press.
- Howes, David, and Constance Classen. 1991. Conclusion: Sounding Sensory Profiles. In *The Varieties of Sensory Experience*, edited by D. Howes. Toronto: University of Toronto Press.
- Iggo, A. 1982. Cutaneous sensory mechanisms. In *The Senses*, edited by H. B. Barlow and J. D. Mollon. Cambridge: Cambridge University Press.
- Jackson, Frank. 1977. *Perception: a Representative Theory*. Cambridge: Cambridge University Press.
- Jeannerod, Marc. 1986. The formation of finger grip during prehension: a cortically mediated visuomotor pattern. *Behavioural Brain Research* 19: 99-116.
- Jeannerod, Marc. 1997. *The Cognitive Neuroscience of Action*. Oxford: Blackwells.
- Kim, Jaegwon. 1998. *Mind in the Physical World: An Essay on the Mind-Body Problem and Mental Causation*. Cambridge, Mass.: MIT Press.
- Kripke, Saul. 1980. *Naming and Necessity*. Oxford: Blackwells.
- Krueger, Lester E., and David Katz. 1989. *The World of Touch*. Hillsdale, N.J.: L. Erlbaum Associates.
- Lederman, S. J., and R. L. Klatsky. 1987. Hand movements: a window into haptic object recognition. *Cognitive Psychology* 19: 342-368.

- Lederman, S. J., and R. L. Klatsky. 1990. Haptic object classification: knowledge driven exploration. *Cognitive Psychology* 22: 421-59.
- Leon, Mark. 1988. Characterising the senses. *Mind and Language* 3 (4): 243-270.
- Lewis, David K. 1969. *Convention: a Philosophical Study*. Cambridge, Mass: Harvard University Press.
- Lewis, David. 1972. Psychophysical and theoretical identifications. *Australasian Journal of Philosophy* L (3): 249-58.
- Lewis, Gilbert. 1975. *Knowledge of Illness in a Sepik Society: a Study of the Gnau, New Guinea*. London: Athlone Press. University of London.
- Loar, Brian. 1996. Comments on Campbell. In *Perception*, edited by E. Villanueva. Atascadero: Ridgeview.
- Locke, John. 1689. *An Essay Concerning Human Understanding*. Edited by P. H. Niddich. Oxford: Clarendon Press.
- Lowe, E.J. 1992. Experience and its objects. In *The Contents of Experience*, edited by T. Crane. Cambridge: Cambridge University Press.
- Lycan, William G. 1996. *Consciousness and Experience*. Cambridge, Mass.: MIT Press.
- Mackie, J.L. 1977. *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- Malpas, R.M.P. 1968. The Location of Sound. In *Analytical philosophy: Second series*, edited by R. J. Butler. Oxford: Blackwell.
- Marr, David. 1982. *Vision*. New York: W. H. Freeman.
- Martin, M.G.F. 1992. Sight and Touch. In *The Contents of Experience*, edited by T. Crane. Cambridge: Cambridge University Press.
- Martin, M.G.F. 1995. Bodily awareness: a sense of ownership. In *The Body and The Self*, edited by J. L. Bermudez, A. Marcel and N. Eilan. Cambridge, Mass.: MIT Press.
- Martin, M.G.F. 1997. The Shallows of the Mind. *Proceedings of the Aristotelian Society, Supplementary Volume LXXI*: 75-98.
- Martin, M.G.F. 1998a. Bodily sensations. In *Routledge Encyclopedia of Philosophy*, edited by E. Craig. London: Routledge.
- Martin, M.G.F. 1998b. Setting things before the mind. In *Current Issues in the Philosophy of Mind*, edited by A. O'Hear. Cambridge: Cambridge University Press.
- Martin, M.G.F. forthcoming. *Uncovering Appearances*.

- McCarthy, Rosaleen. 1993. Assembling routines and addressing representations: An alternative conceptualisation of 'what' and 'where' in the human brain. In *Spatial Representation*, edited by N. Eilan, R. McCarthy and B. Brewer. Oxford: Blackwell.
- McCulloch, Gregory. 1993. The very idea of the phenomenological. *Proceedings of the Aristotelian Society* XCIII: 39-57.
- McGinn, Colin. 1988. Content and consciousness. In *The Problem of Consciousness*. Oxford: Blackwells.
- McGinn, Colin. 1997. *The Character of Mind*. 2nd ed. Oxford: Oxford University Press.
- Massaro, Dominic W. 1998. *Perceiving Talking Faces: From Speech Perception to a Behavioural Principal*. Cambridge, Mass.: MIT Press.
- Mellor, D H. 1977. Natural kinds. *British Journal for the Philosophy of Science* 28: 299-312.
- Michotte, A. 1963. *The Perception of Causality*. Translated by Heath, P. New York: Basic Books.
- Mill, John Stuart. 1979. *An Examination of Sir William Hamilton's Philosophy, and of the Principal Philosophical Questions Discussed in his Critings, Collected Works of John Stuart Mill; Vol. 9*. Toronto: University of Toronto Press.
- Millar, Susanna. 1997. *Reading by Touch*. London: Routledge.
- Millikan, Ruth Garrett. 1991. Perceptual content and Fregean myth. *Mind* C (4): 439-459.
- Millikan, Ruth Garrett. 1997. Images of identity: in search of modes of presentation. *Mind* 106 (423): 499-519.
- Milner, A. David , and Melvyn A. Goodale. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.
- Mishkin, M., L.G. Ungerleider, and K.A. Macko. 1983. Object vision and spatial vision: Two cortical pathways. *Trends in Neurosciences* 6: 414-417.
- Modrak, Deborah K. W. 1987. *Aristotle: the Power of Perception*. Chicago: University of Chicago Press.
- Moore, G.E. 1918-19. Some judgements of perception. *Proceedings of the Aristotelian Society* XIX.
- Moore, G.E. 1957. Visual sense data. In *British Philosophy in the Mid-Century*, edited by C. A. Mace. London: George Allen & Unwin Ltd.

- Morgan, Michael J. 1977. *Molyneux's Question: Vision, Touch, and the Philosophy of Perception*. Cambridge: Cambridge University Press.
- Nakayama, Ken, Zijiang J He, and Shinsuke Shimojo. 1995. Visual surface representation: a critical link between lower-level and higher-level vision. In *Visual Cognition Volume 2*, edited by S. M. Kosslyn and D. N. Osherson. Cambridge, Mass.: MIT Press.
- Nelkin, Norton. 1990. Categorising the senses. *Mind and Language* 5 (2): 149-165.
- O'Neill, Daniella K., Janet Astington, and John H. Flavell. 1992. Young children's understanding of the role that sensory experiences play in knowledge acquisition. *Child Development* 63: 474-490.
- O'Shaughnessy, Brian. 1985. Seeing the Light. *Proceedings of the Aristotelian Society* 85: 193-218.
- Peacocke, Christopher. 1982. *Sense and Content*. Oxford: Clarendon Press.
- Peacocke, Christopher. 1985. Consciousness and other minds. *Proceedings of the Aristotelian Society Supplementary Volume* 59: 97-117.
- Peacocke, Christopher. 1985b. Imagination, Experience and Possibility. In *Essays on Berkeley*, edited by H. Robinson. Oxford: Clarendon Press.
- Peacocke, Christopher. 1986. *Thoughts: An Essay on Content*. Oxford: Basil Blackwell.
- Peacocke, Christopher. 1989. Perceptual content. In *Themes from Kaplan*, edited by H. K. Wettstein, D. Kaplan, J. Almog and J. Perry. New York: Oxford University Press.
- Peacocke, Christopher. 1992. *A Study of Concepts*. Cambridge, Mass.: MIT Press.
- Peacocke, Christopher. 1999. *Being Known*. New York: Oxford University Press.
- Pears, David. 1972. *What is Knowledge?, Essays in Philosophy*. London: Allen and Unwin.
- Perkins, Moreland. 1983. *Sensing the World*. Indianapolis: Hackett.
- Price, H. H. 1932. *Perception*. London: Methuen.
- Proust, Marcel. 1925. *The Guermantes Way*. Translated by Scott-Moncrieff, C. K. London: Chatto and Windus.
- Putnam, Hilary. 1967. The nature of mental states. In *Mind Language and Reality: Philosophical Papers Volume 2*. Cambridge: Cambridge University Press.
- Putnam, Hilary. 1975. The meaning of 'meaning'. In *Mind Language and Reality: Philosophical Papers Volume 2*. Cambridge: Cambridge University Press.

- Rée, Jonathan. 1999. *I See a Voice*. London: Harper Collins Publishers.
- Reid, Thomas. 1983. An Inquiry into the Human Mind and the Principles of Common Sense. In *Thomas Reid: Inquiry and Essays*, edited by R. E. Beanblossom and K. Lehrer. Indianapolis: Hackett.
- Ritchie, Ian. 1991. Fusion of the faculties: a study of the language of the senses in Hausaland. In *The Varieties of Sensory Experience*, edited by D. Howes. Toronto: University of Toronto Press.
- Rivlin, Robert, and Karen Gravelle. 1984. *Deciphering the senses : the expanding world of human perception*. New York: Simon and Schuster.
- Robinson, Howard. 1994. *Perception*. London: Routledge.
- Rock, Irvin. 1975. *An Introduction to Perception*. New York: Macmillan.
- Roessler, Johannes. 1999. Perception, introspection and attention. *European Journal of Philosophy* 7 (1): 47-64.
- Roxbee-Cox, J. W. 1970. Distinguishing the senses. *Mind* 79: 530-550.
- Russell, Bertrand. 1912. *The Problems of Philosophy*. Oxford: Oxford University Press.
- Sanford, David H. 1976. The primary objects of perception. *Mind* LXXXV.
- Schiller, P. H., and N. K. Logothetis. 1990. The colour-opponent and broad-band channels of the primate visual system. *Trends in Neuroscience* 13: 392-98.
- Shallice, Tim. 1988. *From Neuropsychology to Mental Structure*. Cambridge: Cambridge University Press.
- Smith, A.D. 1990. Of primary and secondary qualities. *The Philosophical Review* 99: 221-251.
- Sorabji, R. 1971. Aristotle on demarcating the five senses. *Philosophical Review* 80: 55-79.
- Squire. 1989. Mechanisms of memory. In *Molecules to Models: Advances in Neuroscience*, edited by K. L. Kelner and D. E. Koshland. Washington: American Association for the Advancement of Science.
- Stein, Barry E., and M. Alex Meredith. 1993. *The Merging of the Senses*. Cambridge, Mass.: MIT Press.
- Strawson. 1959. *Individuals*. London: Methuen.
- Strawson, P.F. 1979. Perception and its objects. In *Perceptual Knowledge*, edited by J. Dancy. Oxford: Oxford University Press.

- Theophrastus. 1987. On the Senses. In *Early Greek Philosophy*, edited by J. Barnes. Harmondsworth: Penguin.
- Tye, Michael. 1992. Visual qualia and visual content. In *The Contents of Experience*, edited by T. Crane. Cambridge: Cambridge University Press.
- Urmson, J.O. 1968. The objects of the five senses. *The Proceedings of the British Academy*.
- Warren, D. H. , R. B. Welch, and T. J. McCarthy. 1981. The role of visual-auditory 'compellingness' in the ventriloquism effect: Implications for transitivity among the spatial senses. *Perception and Psychophysics* 30: 557-64.
- Welch, R.B., and D.H. Warren. 1978. Intersensory interactions. In *Handbook of Perception and Human Performance*, edited by K. Bluff and L. Kaufman.
- Westfall, Richard S. 1980. *Never at Rest. A Biography of Isaac Newton*. Cambridge: Cambridge University Press.
- Williams, Bernard. 1978. *Descartes: the Project of Pure Enquiry*. Harmondsworth: Penguin Books.
- Zeki, Semir. 1993. *A Vision of the Brain*. Oxford: Blackwell Scientific.

Acknowledgments

I would like to thank my supervisor, Mike Martin, for his help