

The Henry Hardy Virtual Library

Subjective Experience

Henry Hardy

This is the text of the thesis submitted by the author for the degree of DPhil to the Faculty of Literae Humaniores in the University of Oxford in Hilary Term 1976. No changes of substance have been made, but some subeditorial tidying has been done.

Footnotes in square brackets are comments made by the author after the award of the degree. In the case of the first chapter, some of these are or derive from notes made as part of an attempt to prepare the chapter for publication, a project that was soon abandoned. Bold red arabic numbers in square brackets mark the beginning of the relevant page in the original typescript. Cross-references are linked to the pages of the present text.

The thesis is in effect a second edition of the author's BPhil thesis, '**Subjective Experiences**'.

A concise review of both theses by Johnny Lyons is available **on this site**.

© Henry Hardy 1976, 2018

First Posted in the Henry Hardy Virtual Library 30 November 2018
Reformatted 4 April 2021

Preface

I should like to express my thanks to my supervisors, Anthony Kenny and David Pears, for their indispensable help and advice; and also to Samuel Guttenplan for useful discussion, and for much help with the practical side of preparing a thesis.

HRDH
Wolfson College
April 1976

Note: This thesis is submitted under the new regulation of the Lit. Hum. board, whereby holders of the BPhil are permitted to submit a shorter thesis (of not less than 45,000 words) for the degree of DPhil. This thesis, while on the same subject as my BPhil thesis, is with the exception of a few short passages entirely new.

Contents

Preface	iii
Introduction	1
Abstract	5
1 <i>Sensations as designated objects</i>	10
The ‘object and designation’ model of sensation language	12
Objections to ODM	15
The argument from analogy	15
‘Here is S again’: argument type 1	21
‘Here is S again’: argument type 2	31
(a) Recognition of a sensation as of type S	45
(b) Calling to mind what S is like	45
(c) Remembering what ‘S’ means	46
What Wittgenstein said	51
Intrasubjectivity and intersubjectivity	56
2 <i>C-subtle</i>	61
Arguments against C-subtle	66
3 <i>What do sensation terms mean?</i>	88
A framework of enquiry	91
Methods of enquiry	95
Teaching links	97
Changes of nature: pain	110
A general hypothesis	130
Conclusion	143
Bibliography	147

INTRODUCTION

[Men] *suppose their words to be marks of the ideas in the minds also of other men, with whom they communicate.* for else they should talk in vain, and could not be understood, if the sounds they applied to one idea were such as by the hearer were applied to another, which is to speak two languages.

Locke (1690) 3. 2. 4 (his italics)

This thesis aims to defend a simple but important belief: that a crucial part of communication between human beings is communication about subjective experience. Such communication often takes place directly, when people talk to one another about their feelings – their sensations, emotions and moods. But also the assumption of shared subjective experience in shared circumstances lies behind the whole enterprise of linguistic intercourse, in such a way that if this assumption could be shown to be false, and if, further, it became generally known and understood that it was false, not only would there be great distress, but even if communication could continue at all, it would be only in some heavily attenuated form as compared with what we now believe ourselves capable of, and we would have to establish a very different set of attitudes towards it.

The belief I am defending is one that is taken for granted at a common sense level. The need to defend it arises only because philosophical objections have been lodged against it, and have secured a wide following. [2] According to the objectors, the fact that we can never achieve anything so direct as introspection of the subjective experience of any person other than ourselves means that the whole picture of people as communicating about or in terms of their subjective experience is, in some more or less

INTRODUCTION

radical way, misconceived. Though our ordinary ways of speaking may not in themselves require amendment, they must not be misunderstood or misinterpreted according to a simple-minded model of the relation between language and experience, whereby experience is thought of as a realm of objects that can be straightforwardly designated by linguistic terms just as ordinary material objects can. As has often been said, we cannot necessarily talk sense about concepts simply in virtue of being able to talk sense with them.¹ And mental concepts in particular are notoriously susceptible (so it is claimed) to being misunderstood by the philosophically naive.

In my view, exactly such a 'simple-minded' model is the correct one; and the difficulties supposed to lie in the way of its application are (where they are real) not insuperable. According to this model, subjective experiences are private to the subject. He alone is directly aware of their occurrence and of their phenomenal nature. There is no possibility at present of one subject of experience comparing his experiences with those of another subject [3] to discover immediately whether or not their experiences are qualitatively similar in similar circumstances. There is no way for a child who is learning language to reveal the quality of his experiences to the adults who are teaching him, so that they may be able, when telling him the names of his various subjective experiences, to have the accuracy of their pedagogical labelling confirmed by direct inspection of the quality of the items being labelled. There are, of course, indirect means of doing all these things, means that we constantly use. They are adequate for their purpose, but they still leave room for doubt, and by the same token for faith, the faith that we all subscribe to, in our subjective experiences being interpersonally similar in range and disposition of quality. My inner world, we believe, maps on to our shared public world in much the same way as does yours.

¹ E.g. Ryle (1949), introduction.

INTRODUCTION

This thesis is taken up largely with a defence of this model, and so the discussion is mainly about the nature of sensation language. A given sensation term, I argue, can, does and must refer to experiences of the same subjective quality for different people, as well as to sensations whose public manifestations are the same. Only in this way is interpersonal communication about subjective experience possible.

But the implications of the model of sensation language I defend are wider. The admitted difficulties of discussing subjective experience at all, let alone doing so in precise [4] terms, has created blinkers, and not only methodological ones, in academic disciplines other than philosophy. Most human sciences – sociology, psychology, social anthropology and the rest – fight shy of giving subjective experience due weight in their deliberations.² Even ordinary natural science is normally conducted as if the world could be explained without reference to the subjects who comprehend it. This is understandable, since methodologies which exclude subjective experience are more productive, in terms of sheer quantity of accurate, detailed and useful information. It may therefore be reasonable to adopt such methodologies for certain purposes. But they must not become imperceptibly metamorphosed into ontologies. In the last analysis, [5] subjective experience will have to have its proper place in our understanding of the world, and no ossified methodology, adopted originally for perhaps acceptable reasons

² Things are beginning to look up. In a recently published introduction to psychology we read: 'It might be questioned whether there is [...] any point in being concerned about private experience, at least from a scientific point of view. There is a point, because even though hypotheses [about similarities and differences between different people's sensations] cannot be tested, that does not make them *ipso facto* false.' Legge (1975), 30. Cf. Laing (1971), 21: 'People may be observed to sleep, eat, walk, talk, etc. in relatively predictable ways. We must not be content with observation of this kind alone. Observation of behaviour must be extended by inference to attributions about experience. Only when we begin to do this can we really construct the experiential-behavioural system that is the human species.'

INTRODUCTION

of expediency, should be used as a basis for shirking this prospect.

But these considerations, important though they are, lie beyond the scope of the main argument of this thesis, which merely defends the model of our thought and talk about subjective experience whose demise would make genuine discussion of this side of life impossible. This is only a preliminary task, as so often in philosophy, but it is not on that account less important: if anything it is more so, since so many crucial possibilities depend on its successful execution. [6]

ABSTRACT

Chapter 1

In the first chapter I state the view, and defend its possibility, that sensations are like ordinary objects inasmuch as the language which we use to talk about sensations contains terms which designate the sensations in much the same way as names of ordinary objects designate these objects: that sensations are identified both by their intrinsic subjective quality and by the public surroundings in which they occur; that these two kinds of criterion are jointly necessary for identifying sensations; and that they apply interpersonally, so that a sensation of a given type must have the same intrinsic quality no matter whose sensation it is.

A number of arguments against the possibility of this view are considered. They all impinge on the well-known ‘argument from analogy’, whose intelligibility and validity are essential to the view I defend. The first objection to the argument from analogy which I consider attacks its validity: it is claimed that it is a weak argument. I reply that it is stronger than is sometimes thought, and that even if it still falls short of some ideal standard, that is no reason to reject it. Its acceptance is deeply embedded in our way of thinking of other people, and we would need very persuasive arguments to show that it must be [7] rejected.

Next I consider arguments against the premiss of the argument from analogy – the premiss that I have subjectively similar experiences in objectively similar circumstances. It is claimed either that this premiss, depending as it does on the veracity of memory, is too insecurely supported to be relied on; or that it is actually unintelligible if thought of in terms of the

ABSTRACT

view of sensation language I adopt. Against the first kind of argument I maintain that it is not unreasonable or irresponsible to rely on the accuracy of memory; against the second, that the verificationist conditions of meaningfulness on which it is based are irrelevantly severe. This involves me in a defence of the view that the conditions of meaningfulness in question are indeed verificationist; and in demonstrating that Wittgenstein, who put forward an argument of this second kind, was a straightforward verificationist sceptic about memory, rather than something more subtle.

Finally, having thus defended the intelligibility and truth of the premiss of the argument from analogy, I defend the intelligibility of its conclusion – that different people have subjectively comparable experiences in comparable circumstances. I do this by discussing various conceivable ways in which the subjective quality of the experience of two people could be directly compared. Given that the conclusion of the argument from analogy is intelligible, and that my defence of the validity of the argument is successful, it now follows from the [8] truth of its premiss that the conclusion too is true. Hence, the view I adopt at the beginning of the chapter is possible, at least. If it is shown to be false, it will not be on a priori grounds.

Chapter 2

In chapter 2 I show how an alternative analysis of sensation language, which appears to be immune to some of the criticisms levelled against the analysis I adopt in chapter 1, and also avoids the behaviourist extreme, is not as real a possibility as it at first seems. Its initial plausibility depends on neglect of its more far-reaching implications.

According to this theory ('C-subtle'), sensations are defined exclusively in terms of their public surroundings, not by their intrinsic quality – though this must be constant for any given person (except on an extreme interpretation of the theory, which I quickly reject).

ABSTRACT

I give three main arguments against this theory. The first argument begins by pointing out that, according to the theory, it should be possible for someone to know the meaning of any sensation term without ever having experienced the sensation in question, so that he might fail to recognise the sensation when he first felt it, despite his knowledge of the meaning of its name. I argue that this paradoxical situation is unacceptable for more general sensation terms such as ‘pain’, even if it is tolerable for more specific ones such as ‘headache’. This distinction [9] is the occasion for some general remarks about how much experience is required for knowing the meanings of words.

The second argument, which arises out of the first, is that C-subtle’s requirement of intrasubjective constancy of sensation quality for a given type of sensation, together with some conclusions of Strawson’s, entails precisely the requirement of intersubjective similarity which C-subtle wishes to eschew.

The third argument is that the consequences of applying C-subtle, not just to individual sensation terms, but to the term ‘sensation’ itself, are unacceptable. There are two main consequences in question. One is that we could not even require that our sensations are in general like those of other people. And the other consequence, which would be a sufficient objection to C-subtle even if the first consequence did not matter, is that the term ‘sensation’ cannot be made intelligible at all on C-subtle’s view of its meaning, since there is no more general public concept, already mastered, under which ‘sensation’ could be subsumed and in terms of which the meaning of ‘sensation’ could be understood. We could not learn what sort of thing a sensation was unless, contrary to C-subtle, we allowed ourselves to be instructed by our own personal experience.

Chapter 3

The last chapter consists of a preliminary discussion of the question of the actual – no longer merely the possible – [10] meanings of sensation terms. There is a range of possible views

ABSTRACT

as to what sensation terms mean, from the ‘intrinsicist’ extreme of holding that they are defined purely in terms of inner quality, to the ‘extrinsicist’ extreme of holding that they are defined exclusively in terms of their public surroundings. Between these extremes lies the middle possibility that both intrinsic and extrinsic criteria are involved.

These are the broad options. Finer discrimination is possible, and I try next to sketch a framework within which a more detailed typology of sensation terms could be established. I subdivide the categories both of extrinsic and of intrinsic criteria, and suggest some questions that arise out of the resulting classificatory scheme: questions as to what general truths about the meanings of sensation terms might emerge, and questions about the explanation of these truths, once discovered.

Next I outline, and demonstrate in action, two possible methods for investigating the meaning of a sensation term, both of which can help to reveal not only that element in the meaning of a term which is closely related to its functional role in our lives, but also those elements in its meaning which arise from the associative power of what, from a functional viewpoint, are inessentials. The first method is to inspect the ‘teaching links’ of a sensation – the public phenomena by reference to which the use of the term is taught. I discuss what the various teaching links of sensation terms are, and what their twofold connection with meaning is. I also mention two [11] ways in which they are unreliable guides to meaning: In particular, they can tell us nothing about intrinsic criteria.³

This is where the second method comes in. It is to ask what we would be disposed to say if nature were to change in certain ways. I exemplify this method in connection with ‘pain’: changes of intrinsic quality and of extrinsic surroundings are imagined, in order to try to elicit the criteria which govern the use of the term. Heavy reliance is placed, in considering what we should say if such changes occurred, upon the deliverances of intuition: this, I

³ [And links are not necessarily implicated in meaning.]

ABSTRACT

argue, is the only method available in the last resort for settling the issues in question, though support for its rulings can be supplied, up to a point, by examining the function that a sensation concept performs in our lives – and this is one of the things that the two methodological tools described above are designed to reveal.

Out of this discussion arises a general hypothesis, that all our sensation concepts, being tailored to the world as we know it, apply only when both intrinsic and extrinsic criteria are satisfied, and so only when private sensation quality and public surroundings march in step. Were these two kinds of criterion to part company, our concepts would need either to fall into disuse, or to be given new definitions⁴ that fitted the novel circumstances – in which case they would not really be the same concepts at all. I show this hypothesis in action by outlining a couple of applications of the framework of enquiry drawn up earlier in the chapter. [12]

⁴ [Or uses.]

1 SENSATIONS AS DESIGNATED OBJECTS⁵

What do sensation terms mean? Wittgenstein recommends (1953, §293) that we do not ‘construe the grammar of the expression’⁶ of

⁵ [Or ‘What Are Sensations?’]

⁶ The terms in which Wittgenstein’s recommendation is expressed need adjustment. Wittgenstein talks of ‘the grammar of the *expression* of sensation’ (my italics) rather than simply of ‘the grammar of sensation language’ as might be expected. Why does he do this? Presumably because his own theory of sensation language at this time was that it is learnt by means of the substitution of unanalysable verbal utterances for the natural expressions of sensations (see e.g. §244). So sensation talk is, for him, just a way of expressing sensations, and thus the grammar of sensation language is identical with the grammar of the (verbal) expression of sensation. But what is at issue is precisely whether Wittgenstein’s view of sensation language is correct or not: so he should not have allowed his own theoretical predilections to be embodied in the argument in this way. By doing so, he begs the question. Rival views should be given a proper chance at this stage of the argument, by being discussed as far as possible in theory-neutral terms. Here the rival view is the ‘object and designation’ model of sensation language. Obviously to express a sensation is not simply to designate an object: but it is not at all obvious that we do not designate an object when we use a sensation term in a description or report. And even if ‘I have a pain’ is construed as a mere expression of pain, it seems arguable that a pain is designated. To express is not to designate; but in expressing I may, *inter alia*, designate. Wittgenstein may believe that we cannot use sensation terms in this way: but he must not assume it in advance. True, his reference to expression does not in fact affect his argument in §293. But it is still a misplaced allusion to a rival view, and so worth eliminating. So let us consider his recommendation, that we avoid the ‘object and designation’ model of sensation language, as if it were shorn of its reference to expression.

[13]

Wittgenstein does in some passages allow or imply that we can describe sensations, but he insists that in this case ‘describe’ will be used in a special

SENSATIONS AS DESIGNATED OBJECTS

sensation on the model of “object [13] and designation”): we must not think of sensations as objects, designated by terms in the same way that material objects are designated. If we do think of sensations in this way, Wittgenstein says, ‘the object drops out of consideration as irrelevant’, because it would make no practical difference to sensation talk what the object was like, whether it was the same for different people, or indeed whether there was any object at all. But this would be unacceptable: sensations are not, surely, irrelevant in this way to sensation language. So if this way of looking at sensation terms yields the result that sensations are irrelevant, it must be mistaken. Accordingly, Wittgenstein presumably believes, some model other than that of ‘object and designation’ (hereafter ‘ODM’), a model whereby sensations retain their role in sensation language, must be adopted.

Against this I shall defend a version of ODM as the right model to use in analysing sensation language. Certainly the use of names of sensations is not in all ways comparable to the use of names of ordinary material objects, and the differences between the two should not be ignored. But I believe that, as far as the disagreement between Wittgenstein and myself is concerned, the similarities are more important than the differences. [14]

This first chapter, then, is devoted principally to the defence of the possibility of a version of ODM of which I give an initial specification below (consideration of a rival version of ODM, and more detailed treatment of the version I favour, are reserved for later chapters). Of course, the successful defence of a possibility is not the same as the establishment of an actuality. But this preliminary exercise is necessary because, unless the defence of the possibility of ODM is successful, the option would

sense – not the same as occurs in ‘I describe my room.’ I believe that the two kinds of description, though obviously different, are more closely analogous than Wittgenstein thinks (see, e.g., §§290–1).

not be open, after the rejection of alternatives,⁷ [15] of falling back on this model as the true one.

The ‘object and designation’ model of sensation language

First of all we need a rough characterisation of the version of ODM I am sponsoring. In providing this I follow well-established practice in the philosophical literature, as I shall do again in later pages, by using the example of pain. Concentration on the case of pain can often lead to over-generalisation of the special features⁸ of pain, or to neglect of the special features of

⁷ I do not consider these alternatives in the thesis: indeed I am not at all clear what they are, apart from the one attributed to Wittgenstein on p. 1, which, though not untrue, is only a very small part of the truth. Interpreters of Wittgenstein have attempted in various ways to fill out the smattering of remarks he makes into a theory of sensation language, but the positive characteristics of such a theory are always rather elusive (the notion of ‘necessarily good evidence’ is particularly troublesome: cf. Hacker (1972) and Hardy (1974), pp. 87–94). It is easier to state what the theory rules out. It denies that discussion of the raw phenomenal events in our consciousness makes sense, let alone leads to the answering of questions about intersubjective similarities or differences, or about the inferring of the occurrence of subjective experiences from the observation of objective phenomena. So the term ‘subjective experience’ will not, in so far as it means anything, refer to events or objects private to a given subject: a person’s experiences will only be subjective in the sense that they are his and not someone else’s. Nor will a subject of experience be the only person with a certain kind of privileged awareness of his experiences, except again in the trivial sense that, since the experiences are his, he is bound to be present, and so aware of them, when they occur. Nor will he be a unique authority on the quality of his experiences. The question whether two people have comparable experiences in comparable circumstances will, in so far as it is meaningful, be easy to answer. And the adult teaching the child the language of subjective experience will be indulging in no risky inferences. [15]

Since the need for alternative theories arises only from a rejection of ODM, their omission from the discussion may not be too serious. I do, however, discuss at some length in chapters 2 and 3 which of the many possible versions of ODM is to be preferred.

⁸ It would be instructive, no doubt, to enquire just what are the features of pain in virtue of which it is always in the philosophical limelight.

SENSATIONS AS DESIGNATED OBJECTS

other sensations. These are dangers to be aware of. But for the purposes of this prefatory sketch, at least, pain may serve.

Pain, then, is a private, subjective⁹ experience. Only [16] I am directly aware of my own pain, and only from my own case do I know what pain feels like. If I have never felt pain, the understanding of the meaning of ‘pain’ that I can acquire is severely limited in a way specified below in chapter 2 (pp. 72–75). I discover that others are in pain by noticing its effects, or its causes, or because they tell me. I believe that other people’s pain feels like my own – indeed it is part of the meaning of ‘pain’ that it should do so – though I cannot verify this belief by inspecting their pain and comparing it with my own. By the same token I cannot be as sure that others are in pain as they are themselves: there is always the possibility of deception or error.

The way in which the truth of the belief that the pain of others feels like mine is essential to the meaning of ‘pain’ is this: it is at least one of the necessary conditions for the correct ascription of pain that a sensation be occurring with the right subjective quality. We assume, on good grounds, that this condition is met both in our own case and in the case of others; and this kind of assumption is the basis both of communication about our sensations and of the teaching of the meaning of sensation terms to children learning their first language.

This much said, it must now be added that the subjective quality of a sensation is not the only important criterion that determines whether or not the sensation is a pain. Pain is also essentially something we in general wish and [17] try to avoid; it is essentially associated with certain characteristic behaviour

⁹ All experiences, in the sense of ‘experience’ used here, in which the word refers to an event or episode in the stream of consciousness, are necessarily private and subjective. To this extent it is pleonastic to describe them as such. My purpose in including these epithets is to draw attention to those features of sensations, necessary though they may be, whose existence, real or supposed, provides the main grounds for the objections to ODM considered in this chapter. These features must be explicitly faced up to by anyone who wishes to defend ODM as the true model, or even only as a logically possible one.

SENSATIONS AS DESIGNATED OBJECTS

(wincing, screaming, tending etc.); and it is essentially responded to, again in general, with sympathy and sometimes treatment. Pain, as it were, stands at the centre of a complex structure of cause and effect, of instinctive response and conventional social behaviour, without which the concept of pain would be very different. This kind of insight, which I have characterised here with extreme sketchiness, we owe principally to Wittgenstein; and it is extremely important.

All the same, it is an insight that tends to be overdone. The subjective quality of the sensation can be forgotten, or even disowned, amid the emphasis placed on the other aspects of the concept. The fact that we can know of the pain of others, and teach the use of 'pain' to children, only on the basis of the causes, effects and other typical concomitant circumstances of pain, is used as a reason for belittling the significance of the subjective quality of pain: either it is allowed that pain has such a quality, but denied that this quality plays any role in the meaning of 'pain'; or, worse, it is said that the very notion of a 'subjective quality' as distinct from public manifestations is an empty one. Against this I shall argue that subjective quality is in fact as crucial an element in the structure of the concept as any other.

This account has been deliberately expressed in terms which, according to the opponents of ODM, embody the very [18] false model that causes all the trouble; terms which beg all kinds of questions and perpetuate all kinds of myths. But it is what I believe to be the true view, expressed as straightforwardly as possible. Since I hold that this view is right, and that the terms in which it is expressed are perfectly adequate, standing in need of no revision or qualification, I have expressed it simply, uncluttered by the hedging circumlocution and qualification-ridden defensiveness which so often characterises statements of a similar kind. I have done this because prophylactic measures of this kind prove, when the argument is finished, to be unnecessary. None of this, of course, counts as an excuse for dodging the many criticisms levelled against ODM: indeed it is these which I shall now consider. But it is worth stating one's

SENSATIONS AS DESIGNATED OBJECTS

initial position forthrightly, without the implicit concession of compromising one's mode of expression in the light of arguments whose force one eventually wishes to deny.

Objections to ODM

There is more than one kind of objection to the view I have sketched. There are arguments that the view is (intelligible but) false because the model it espouses is based on an unjustified reliance on memory, or on unprovable assumptions about interpersonal similarities. And more radically there are attempts to show that the terms in which the view is stated are such as to make it unintelligible: [19] to show that our 'understanding' of what is being said is an illusion. The upshot of this latter approach would be that, if ODM were true, sensation terms would be meaningless: and so, it is argued, ODM must be rejected.

There has not been agreement, as will appear below, about which type of objection Wittgenstein adopts in his so-called 'private language argument'. But it is enough that both kinds of objection have been advanced, and I will now explain why I believe them to be of insufficient force to rule out ODM.

The argument from analogy

I conclude that other human beings have feelings like me, because, first, they have bodies like me, which I know, in my own case, to be the antecedent condition of feelings; and because, secondly, they exhibit the acts, and other outward signs, which in my own case I know by experience to be caused by feelings. I am conscious in myself of a series of facts connected by an uniform sequence, of which the beginning is modifications of my body, the middle is feelings, the end is outward demeanour. In the case of other human beings I have the evidence of my senses for the first and last links of the series, but not for the intermediate link. I find, however, that the sequence between the first and last is as regular and constant in those other cases as it is in mine. In my own case I know that the first

SENSATIONS AS DESIGNATED OBJECTS

link produces the last through the intermediate link, and could not produce it without. Experience, therefore, obliges me to conclude that there must be an intermediate link; which must either be the same in others as in myself, or a different one: I must either believe them to be alive, or to be automatons: and by believing them to be alive, that is, by supposing the link to be of the same nature as in the case of which I have experience, and which is in all other respects similar, I bring other human beings, as phenomena, under the same generalisations which I know by experience to be the true theory of my own existence.

Mill (1889), 243–4

[20] An essential element in my version of ODM is the assumption that like sensation quality goes with like behaviour, like stimuli and so forth. I feel subjectively thus-and-so when such-and-such, and I believe that others feel the same way in the same predicament: not just that they would respond in a parallel way, and say the same about how they feel, but that the actual inner quality of their experience is similar to mine. It is this kind of belief – which, I hold, lies at the root of interpersonal communication – that the argument from analogy is an attempt to justify. I say ‘justify’ rather than ‘explain’ deliberately: the argument is a defence against its critics of a belief we already hold, not a hypothesis of developmental psychology about how this belief is acquired. It may well be implausible to suppose that, as infants, we begin by regarding people as unfeeling automata, and only later come to invest them with experiences like our own, when the force of the argument from analogy dawns on us. Such a misconstruction of the role of the argument is represented in figure 1. But the argument is not meant to be a reconstruction of the development through time of a child’s perception of others. Rather it is a retrospective rationalisation, a hindsightful justification, of beliefs we already hold, perhaps instinctively or even innately. It says that the correspondence we find in ourselves between certain qualities of experience and certain kinds of publicly perceptible phenomena is sufficient [22]

evidence for our belief that a correspondence, and a similar correspondence too, also obtains in the case of others. [21]

The argument from analogy is often said to be too weak. Perhaps this is the point Wittgenstein is making when he says (§293) ‘how can I generalise the one case so irresponsibly?’ Inductions, it is held, must be based on more than one instance.

It is also argued, in verificationist mood, that since it is impossible to test the truth of beliefs which rely for their justification on the argument from analogy, such beliefs are at best too insecure to be of crucial importance in the functioning of some of our most basic concepts, and at worst simply unintelligible. If the former of these two views is adopted, the beliefs in question will be discounted, shown to have no role in the correct analysis of mental concepts, though it will not be contended that the terms in which the beliefs are expressed have been improperly understood.¹⁰ But if, as more often happens, the more severe view is adopted, it will be maintained that it makes no sense to interpret the supposed beliefs in such a way that they cannot, even in principle, be justified. In fact (the argument goes) the apparent impossibility of justifying what are undoubtedly important beliefs shows precisely that these [23] beliefs have been radically misconstrued. We must correctly understand the ways of talking which generate the difficulty, and then the difficulty will turn out to have been illusory. Locutions such as ‘the way pain feels to me’, ‘I know just how you feel’, ‘Do you feel the same way as me when you experience such-and-such?’ and the like, if properly understood, should create no mystery: thought about them should not leave us with the sense that there is something inscrutable occurring which, despite its importance to us, we cannot communicate about; something which, in the case of others, we would often dearly like to get closer to, though because of the stubborn facts of nature we are debarred from confronting it face to face. Rather, as Wittgenstein said, any apparent mystery results from some ‘grammatical’ delusion. We

¹⁰ ‘C-subtle’, one of the fruits of this approach, is the subject of chapter 2.

SENSATIONS AS DESIGNATED OBJECTS

are misled by the form of the sentences in which we talk about sensations into conceiving of sensations in a deeply mistaken way.




Fig. 1 *The argument from analogy misconstrued*

The answer to the charge that induction from one case is too insecure to be responsible is, to start with, that it cannot be irresponsible to do one's best. There are areas of experience, perhaps, where better than analogical evidence is available. In

SENSATIONS AS DESIGNATED OBJECTS

these areas, then, let such evidence be made use of. But where the quality of other people's sensations is concerned, analogy is (at present) the only device at our disposal, and for better or worse we must be content with it. What *is* irresponsible is to be so afraid of the 'huge' analogical leap involved in subscribing [24] to the kind of belief in question that one runs into the arms of a theory of sensation language which, while it matches up better to standards of evidence extrapolated from the material object department of language, may fail to give us a true picture of the way in which sensation language functions.¹¹

In any case, the inductive move in the argument from analogy is not as flimsily supported as is sometimes made out in hostile caricatures. I say this for a number of reasons. For one thing, the case which provides the basis for the induction is 'one' case in only one sense: viz., it involves appeal to the experience of only one person. But within this limit there are, in another sense, many cases: [25] each of my subjective types of experience is constantly correlated with its own set of public phenomena, and each such constant correlation provides independent corroboration for the hypothesis that, in general, sensation quality and public circumstances are correlated.

Besides, there is a wealth of evidence, quite outside the realm of sensations, that human beings are similarly constructed and function in similar ways. Indeed, this is so well established that

¹¹ It is important to keep two kinds of question separate: (a) What is the actual structure of our concepts, and what beliefs and assumptions underlie this structure? (b) Should the structure of our concepts be as it is?

It is possible, of course, that the beliefs and assumptions which underlie our current sensation language can be shown to be false. In this case, if the falsity was to be eradicated; at least a radical change of attitude to our communication about sensations would be required – if not an actual restructuring of language.

I believe both that our sensation language does function according to my version of ODM, and that there is no good reason why it should not be allowed to continue doing so. I defend both beliefs in the thesis, but the current discussion of the argument from analogy is mainly concerned with the second one.

SENSATIONS AS DESIGNATED OBJECTS

talk of ‘evidence’ seems unnecessarily cautious: more plausibly, we are concerned with fact. Our physical construction, our behaviour in response to our environment, our internal neurophysiological life histories, are broadly parallel and predictable. Human beings share a common ancestry, and have evolved as a group, facing the same pressures and constraints, and developing common characteristics to serve common needs. Is it then irresponsible to extrapolate from this welter of known homogeneity and hold that sensation quality behaves in a similarly uniform way? If not, here is another source of support for the argument from analogy, itself admittedly analogical, but far from being based, in any sense, on only one instance.

Even if it is allowed that appeal to this richer kind of analogy adds strength to the traditional, self-effacing version of the argument from analogy, it may be said that it does not add enough strength: there is still the verificationist’s insistence that, since a check on the truth of a belief which can be justified only by the [26] argument from analogy is in principle impossible, let alone conclusive, then that belief must be abandoned, or at least radically reinterpreted.

Why should this insistence be yielded to? Well, once again, certainly many of our beliefs can be tested to a higher standard of verification than can be achieved in the case of our beliefs about the interpersonal comparability of sensation quality. But this is not sufficient to show that these beliefs about sensations are ill-founded, mistaken or senseless. All it shows is that we entertain some beliefs on less certain foundations than others. And this is hardly surprising, or indeed irresponsible. On the contrary, it would be impractically idealistic to hold that, if a model of a department of language entails the attribution to us of beliefs which cannot be conclusively checked, then the model must be misconceived.

Besides, the verificationist view in question is dogmatically asserted rather than being argued for: there are, indeed, no checks, let alone conclusive ones, of *its* truth. In which case, perhaps it is sufficient at this stage just dogmatically to assert the

contrary view. There are, of course, considerations which can be urged in defence of the non-verificationist view, and some of these will emerge in due course below.¹² By no means the least of these [27] considerations is that it accords with our intuitions: let this, for the time being, justify at least our not eliminating it a priori.

'Here is S again' – argument type 1

More radically, it is argued that, if ODM is true, not even the *premiss*¹³ of the argument from analogy can be established, let alone its conclusion. The premiss in question is that I, at least, have sensations of a similar subjective quality in similar circumstances. And if this premiss cannot be established, then we cannot even get started on defending the belief, crucial to ODM, that different people have subjectively similar sensations in similar circumstances.

The argument that this premiss cannot be established appears in both of the forms distinguished above on p. 15. It appears as scepticism about the reliability of memory, and it appears as a thesis about the necessary preconditions for meaningfulness. It seems to me that the second form of the argument collapses into the first, for reasons which I shall mention below: but to begin with I shall treat the two as distinct, starting now with the first, which goes as follows. [28]

¹² E.g. 'Intrasubjectivity and intersubjectivity', pp. 40–43 below, where I argue that the supposedly impossible checks could conceivably be carried out. That they could be carried out is of course no guarantee that they would confirm the conclusion of the argument from analogy: but the verificationist objections to the argument concern the [27] availability of checking procedures rather than casting doubt on what the upshot of such procedures might be. More support for the non-verificationist view is to be found in chapter 3.

¹³ Strictly, this is only one of the premisses – the other being that other people are objectively similar to me. But this premiss is not in dispute, so perhaps the favouritism shown to the other premiss may be forgiven. (Is there even a third premiss, that objective similarities betoken subjective similarities?).

SENSATIONS AS DESIGNATED OBJECTS

When a person has a sensation, how is he to tell whether or not this sensation is similar in subjective quality to a sensation he had on a previous occasion? Unless he accepts the force of the argument from the general regularity and patternedness of human life,¹⁴ and thus holds his impression that the sensation is, or is not, of similar quality to be supported by the similarity, or difference, of the circumstances in which it occurs, then he will have to rely on his memory. But memory, as we know, is fallible. So why should it be trusted?

Memory is normally subject to checks. If I think I remember something, I can find out whether I am right by asking, looking, calculating and so forth. But the peculiar thing about remembering the subjective quality of a sensation is that no direct, independent check is available. I am the only authority. And this means, it is argued, that such rememberings should not be trusted.

This line of argument is not self-evidently persuasive. Perhaps in general my memory may need public vindication: but once certified as reliable in the public arena, can it not take on its own private assignments without further instance-by-instance supervision? [29]

Even the claim that memory needs *in general* to be measured against a public yardstick should not be allowed to pass without critical comment. It may seem at first sight that memories of public events and memories of private events are so very differently placed in respect of reliable checkability that the former have better credentials than the latter can ever claim (even though they may, as I have suggested, achieve a degree of respectability parasitically on their more securely supported public counterparts). But when we look more closely, we find, as I shall

¹⁴ As I have said above, I think this argument does have force. But to use it at this stage would be to beg the question, since what is at issue is whether a more direct method of establishing a particular kind of regularity is reliable independently of any such general license to assume regularities. The question is whether memory is a self-sufficient source of information about comparisons over time.

SENSATIONS AS DESIGNATED OBJECTS

explain in a moment, that there is not such an enormous difference between the two kinds of memories after all. In which case it will no longer be so easy to maintain that reliance on memories of private events is legitimised only by appeal to the demonstrably superior reliability of memories of public events. And so it may be possible to concede that, if memory is reliable at all, it can be trusted in the private realm just as well as in the public realm. If so, we shall not need to feel at all shifty when we incorporate faith in the veracity of private memories into our theory of sensation language.

Consider the difference between the private and the public case. In the private case, it is said, we have no means other than our notoriously fallible memories for testing our impressions as to similarities or differences of quality between present and past sensations. In the public case, however, abundant independent evidence is to hand. Well, certainly the evidence is more impressive in [30] one way, involving as it does such a cloud of witnesses: but logically it is deficient in a parallel way to the evidence of private memory. There is no independent test whereby we can verify the collective impression of everybody that some current event bears a relation of similarity or difference to some previous one, and so is appropriately described in such and such a way.¹⁵ We might all be wrong in concert, not only about the immediate point at issue, but also about all the other sorts of thing which we otherwise could have appealed to in order to correct our mistake.¹⁶ The correctness of an impression

¹⁵ As is here implied, similar points can be made if what is at issue is not simply whether a memory is correct, but whether a term is being used with a constant meaning. This comes up in connection with the second kind of criticism of the premiss of the argument from analogy: see pp. 19 ff.

¹⁶ Cf. the hypothesis that God set the world up so as to fool us that it had evolved over an enormously long period: that he placed all the 'evidence' of evolution – fossils and so forth – in the world when he created it. If we accept this, there is no limit to how recently the creation of the world might, for all we know, have occurred. It might have occurred a minute ago, so that we

SENSATIONS AS DESIGNATED OBJECTS

cannot, in the last analysis, be more firmly established than by discovery that everyone shares it.

I do not intend to obliterate the distinction between everyone being under a correct impression and everyone [31] being under an incorrect impression. I am not saying that certain facts about verifiability show that we can properly describe as 'correct' both impressions which are correct and impressions which are not. Of course the only correct impressions are the impressions which are, in fact, correct. The point is only that, although there may be no way to distinguish everyone's being under a correct impression from everyone's being under an incorrect impression, we can still understand the difference between these two radically different states of affairs.

Why then should not a parallel point be made about impressions of private events? These impressions cannot be independently tested, but they are after all correct or incorrect: there is no need for this distinction to be obliterated. I am not suggesting that 'pain' is to be used to refer not only to pain but to any sensation which a person fancies is pain: only that it makes sense to suppose that a person is right or wrong in impressions which cannot be tried before a higher court.¹⁷ We can understand the distinction between a correct and an incorrect impression about the quality of a sensation, however untestable such impressions may be. Indeed it is vital to my argument that such a distinction can and must be preserved. [32]

would have been endowed with all our 'memories' of 'previous experiences' as part of the deception.

It might be said that this hypothesis, since it would make no empirical difference, is empty. Rather, I would say, we cannot so easily dismiss it, for we can certainly understand it: it makes sense. But of course it remains more rational to adopt the usual view. Cf. the hypothesis that everything doubled in size overnight (though there might be some side-effects here).

¹⁷ In the case of 'pain', of course, there are public criteria which can be applied. 'S' would be a better example, but it has not been introduced yet. 'Pain', though, will do, because (if I am right) its having the right subjective quality cannot be tested, even if its having the right public surroundings can.

SENSATIONS AS DESIGNATED OBJECTS

It may be argued that the cases are not parallel; that there is a difference of the following kind. In the case of the memory of sensation quality we can point to a standard of verifiability which is actually in use (viz., the standard applied in the case of memories of public phenomena), and say that the criterion of correctness applied in the private case does not match up to this standard; whereas in the case of memories of public phenomena there is no actual (or conceivable?) superior standard to which we can point. Accordingly, it might be argued, while it would indeed be perverse to deny that impressions about public phenomena could be established as correct or incorrect (for if 'correct' and 'incorrect' were not allowed to apply to impressions about public phenomena, what *would* they apply to?), it is not unreasonable to deny this of private impressions. To which my reply is similar to that used on p. 20 above against another outcrop of verificationism: why accept only one standard of correctness, and that the best available? Better to grant that, though testing the correctness of impressions about private events cannot match up to the same standards as can testing the correctness of impressions about public events, they are, nevertheless, both cases of correctness. There is diversity of standards, but the same correctness. The distinction between a correct impression and an incorrect impression must be made with the best tools available: the difficulty of establishing which of two states of affairs obtains does not show that neither state obtains, let alone that the question as to which obtains is unintelligible. Something is either the case, or not the case. If the tools available for testing private impressions are not as trustworthy as those used for testing impressions about public events, this is not by itself a sufficient reason for eschewing the impressions tested by the inferior method. Why should the world be any respecter of our limited powers of discovery? It continues being thus and so, without reference to us. We must not be debarred from embracing the premiss of the argument from analogy by an addiction to verificationist standards which are not those which actually underlie our sensation talk.

SENSATIONS AS DESIGNATED OBJECTS

Another way of coming to appreciate how far the difference between memories of private and public events can become exaggerated is to imagine cases in which memories of public events happen not to be testable in the ordinary way. Imagine first two cases of checking, one private and one public. The private one is like the case of Wittgenstein writing ‘S’ in his diary to mark the occurrence of a sensation (1953, §258). He put ‘S’ three days ago, and today he has a sensation which he remembers as the same, so he writes ‘S’ again. Beyond this memory there is no criterion of correctness. And the public case is that of a man who remembers a train time by visualising a timetable in his head, and then goes to look at the timetable again (let us suppose it is a poster-timetable on a station hoarding) [34] to check his memory. But when he goes to look at this poster, he finds that a bill-sticker has just covered it with an advertisement. The check which is in principle possible is in this case, as it happens, denied to him. Now surely we will not want to deny that, in this case, if the man’s memory is good, he will be perfectly justified in saying to himself: ‘Oh well, I’m pretty sure I’ve remembered the time right, and that means I’ve ten minutes to spare. I’ll risk going off and buying my ticket before the train goes.’ Then he goes off, buys his ticket, and returns in time to catch the train. The possibility I am suggesting is that instead of looking at the case of the supposed recurrence of a sensation as different from this case because it is in principle uncheckable, we should look at it as comparable to this example of an in principle checkable memory which in a given instance happens to be uncheckable. That is, just as the man on the station is justified in trusting his memory in the circumstances described, so Wittgenstein in the diary example is justified in trusting his memory that he is having the same sensation as he had three days ago. The difference between memories of sensations on the one hand and memories of public objects on the other resides in the logical realm of checkability and not in the mundane realm of reasonable doubt.

This is a point of central importance. Let me bring it out in another way. Take the example of the word ‘blue’. It is often said

of words of this kind that the criteria for their use are public, that someone who denies [35] that delphiniums are blue simply doesn't understand English, and so on. But this should not lead us to deny that once I have learnt what 'blue' means I never again throughout my life, in all probability, need to refer back to the criteria by which I learnt it. I simply remember from day to day what colour blue is. True, I could if I forgot have recourse to others and ask them to remind me about the meaning of 'blue'. But normally I won't need to. Similarly, in the case of Wittgenstein marking the occurrence of a sensation by writing down the letter 'S', once he has decided what sensation is to be called 'S', he probably won't be bothered by the fact that it is impossible for him to effect a public check of the consistency of his use of the term, simply because his memory will serve him well and he won't need to doubt it. The point is that the admitted difference between the checkability of the use of 'blue' and the checkability of the use of 'S' seems to be theoretical rather than practical.

But let us make a concession to theory and suppose that I really do start misremembering what I called S', so that it is of some practical significance to say of my use of 'S' that I only seem to myself to be following a rule whereas I am not really following a rule. The difference is supposed to be that the public language with its public checks is not susceptible to these zany oscillations of meaning. But it is no more absurd to suggest that a similar thing might happen in the case of a public language. Suppose [36] everybody woke up one morning and called everything they had called 'blue' the day before 'red', and vice versa. Let us imagine this happens in the age before writing is invented, in order to avoid the irrelevant complications of checking in books. On this new morning all the public checks of the use of the term 'blue' will confirm the use of the term 'blue' for the colour which was previously called 'red', in precisely the same way as the memory check for the use of the term 'S' in the diary case could conceivably confirm the use of the term 'S' for something which three days ago was not in fact called 'S' at all.

SENSATIONS AS DESIGNATED OBJECTS

Neither the individual in the latter case, nor the public in the former case, has any way of noticing the change. And this is another reason for saying that there is nothing very special about public checking beyond its publicity, beyond the fact that it is safeguarded by the huge unlikelihood of everyone remembering wrong all at once. But it is only a difference of degree which separates this unlikelihood from the unlikelihood of a sensation-diarist's having an inconsistent memory.¹⁸ And when in the realm of logical possibilities, we pay scant regard to likelihoods.

Consider next the case of a babe in a wood, or an infant abandoned on a hillside in a remote and uninhabited region. The child grows to maturity without meeting any other human being, and makes its own way in the world. No shepherd passes by and, having pity on the child, takes it into his own home; no maternal wolf suckles it and becomes its foster parent. [37]

It may be rash to assume that an infant in this predicament would survive. But, granted for the sake of the argument that it does, it would not, presumably, be rash to assume that it makes full use of its memory. It will be trying like any other higher organism to make sense of its environment, to structure the 'blooming, buzzing confusion'¹⁹ of its stream of consciousness in accordance with economical and adaptive hypotheses. The data available to it as it tries (unconsciously, no doubt) to formulate such hypotheses will include both bodily sensations and impressions of the external world received through the various senses – at least, this is what we would call them: whether our imaginary child would distinguish them thus from what we call bodily sensations, or indeed distinguish them at all,²⁰ is part of what is at issue.

¹⁸ [If private checking is no good, nor is public checking.]

¹⁹ William James, *The Principles of Psychology* (London, 1890: Macmillan), vol. 1, p. 488.

²⁰ All impressions are private: some of their causes are public, but there is no such thing as a (literal) public impression. So it is not clear whether, in the absence of a public, the child would need to make the distinction. In our actual predicament, it is convenient to refer to the cause of impressions of public

SENSATIONS AS DESIGNATED OBJECTS

Neither in the case of the sensations, nor in the case of the impressions of public events, is there a public available to provide public checks of the veracity of the child's memory. But, trusting his memory just the [38] same, the child will gradually build up the competence to make sense of, and react productively to, his environment, both outside and inside his skin. Whenever he responds to a stimulus, from without or within, on the basis of its remembered similarity to a previous stimulus, he will be relying on the efficacy of a faculty, his memory, which has never been subjected to public testing. Of course, such testing is still in principle possible for the memories of public events, but because there is no public in our example, this testing is in practice unavailable.

In this imaginary situation, then, memories of public events and memories of private events are on a level. Public checking is ruled out contingently in the first case, necessarily in the second, but for practical purposes these limitations come to the same thing. So if we allow, as we surely do, that this child can, despite his deprived state, use his memory to good effect in respect of his potentially public environment, it seems disingenuous to insist that his competence in the necessarily private realm is any less well grounded. One deprivation could never be removed; the other could: but this discrepancy between counterfactuals has no cash value in the situation envisaged. May we not then be persuaded, once again, that not as much hangs on the distinction between the two sorts of memories as at first seems reasonable? Though memory's functions can in normal circumstances be disproportionately well scrutinised in one of its roles as compared with another, [39] it seems that its reliability is overdetermined in

events, rather than to the impressions in themselves. This separates talk about public events from talk about private events, where there is not such a strong reason to make our vocabulary 'subjective-quality-neutral', so to speak – if only because we don't all feel a tickle simultaneously, as we experience a thunderstorm simultaneously – and where there are reasons for not doing so, viz., that we wish to communicate about subjective experience.

SENSATIONS AS DESIGNATED OBJECTS

its scrutinisable role, rather than underdetermined in its unscrutinisable one.

Ayer (1954) imagines a similar case – that of an infant Robinson Crusoe – in his discussion of the possibility of a private language. It is clear that this character could develop an ‘accidentally private’ language on his own. It will of course be in principle possible to check his use of terms because there are relevant public criteria. But in fact no one else is on the island, so no checking is ever done. To all intents and purposes, therefore, young Crusoe seems to be in the same situation as the diarist with his sensation terms. In one case the language has rules whose use is in principle publicly checkable; in the other case direct public checking is ruled out even in principle. But in fact there is no checking done in either case. So it seems we might claim that to the extent to which we can allow Crusoe to be speaking a language, to the same extent we can allow the diarist to be using meaningful terms. In addition, as Ayer remarks, some of the terms Crusoe invents will refer to his unavoidably private sensations, that is, to the ones without external criteria, so that when Friday comes there will be some words whose use Crusoe cannot teach him. But until Friday does come, the terms Crusoe will be able to teach and the terms he won’t have precisely the same status, and this emphasises even more forcefully the illusoriness of the difference between the two kinds of term. [40]

I have laboured this point somewhat, but resistance to trusting private memories is so deeply embedded in the minds of many philosophers that it is worth making clear how public memories (being after all merely an aggregation of private memories) are no better off than private memories, and how private memories are no worse off than public ones. The upshot is that, until a stronger argument is advanced, there seems no irresistible objection to adopting a model of sensation language whereby memory of the subjective quality of sensations plays a crucial role. It is just such a stronger argument that the second type of argument against the possibility of establishing the premiss of the argument from analogy claims to be. This argument I shall now examine.

'Here is S again' – argument type 2

If we are verificationists, and accept the force of the sort of scepticism about memories of private events described above, then we can use this scepticism as grounds for asserting the meaninglessness of sensation terms as conceived according to ODM. But it has been claimed that, even if we are not verificationists, we can arrive at an identical conclusion by a route which does not depend on any verificationist assumption. This claim, which I believe to be ill-founded, I will now consider in detail.

I shall take as representative of the approach in question an article by Kenny (1971), supplemented by [41] personal conversations and communications. Kenny puts forward his argument as an interpretation of Wittgenstein's argument against the possibility of a private language, a language whose words 'are to refer to what can only be known to the person speaking; to his immediate private sensations' (Wittgenstein 1953, §243). This description fits ODM's view of our sensation language, though Wittgenstein's next remark does not: 'So another person cannot understand the language.'²¹ Our actual sensation language *can* be understood. [42]

²¹ Perhaps Wittgenstein means that, *ex hypothesi*, the constantly accompanying public circumstances which enable us to understand our actual sensation language are, in this imaginary case, absent, or at least sufficiently erratic to provide no basis for communication. If this is so, is he talking not about the sensations that we actually have, but about some science-fictional sensations whose properties, or lack of properties, are irrelevant to the present dispute? Yes and no. To the extent that we do, as things are, talk about sensations, there must be public surroundings that enable us to do this. Equally, though, there are many subtle refinements and distinctions in our sensations which we cannot communicate, because they are not reliably (if at all) correlated with corresponding public subtleties. At least, if we do communicate them, it will be by discovering similarities to or differences from other sensations to which there *is* public access. But even if Wittgenstein is talking to some extent about sensations which differ from the ones we know,

SENSATIONS AS DESIGNATED OBJECTS

Wittgenstein constructs an optimal context in which to propound what amounts to an argument against a certain way (ODM) of viewing our actual sensation language. He considers an extreme case of private reference, where there is not even any attempt to communicate with others. Also, the meaning of a sensation term in his imaginary example is to be determined wholly and exclusively by the subjective quality of the sensation of the person who invents the term: this is not to be just one factor (albeit an important) among others, as I hope to show is the case with our actual sensation language.

Thus the case Wittgenstein considers is an invented one, not an actual one. But exactly the same points he makes about it could, in most cases, also be made about our actual sensation language – about the role, albeit partial, that subjective quality plays in determining the meaning of our actual sensation terms. So the point Wittgenstein has to make is not rendered irrelevant by not being expressed directly in terms of ordinary examples. The function of his hypothetical case, in fact, is precisely to focus his argument by enabling him to concentrate exclusively on that feature of ODM's view of our actual sensation language which is supposedly fatal to ODM.

The extreme case is this: **[43]**

I want to keep a diary about the recurrence of a certain sensation. To this end I associate it with the sign 'S' and write this sign in a calendar for every day on which I have the sensation. – I will remark first of all that a definition of the sign cannot be formulated. But still I can give myself a kind of ostensive definition. – How? Can I point to the sensation? Not in the ordinary sense. But I speak, or write the sign down, and at the same time I concentrate my attention on the sensation – and so, as it were, point to it inwardly. [...] in this way I impress on myself the connection between the sign and the sensation.

Wittgenstein (1953), §258

this does not matter, since, as I say shortly, his argument also applies, in essentials, to our actual sensation language.

SENSATIONS AS DESIGNATED OBJECTS

Wittgenstein maintains that this is an empty procedure.²² There is disagreement about his reasons. Some say he is a sceptic about memory in the way discussed above, and adduce §265 as evidence supporting this interpretation. Others, Kenny among them, say he is making a different sort of point, which I shall now describe. The contention is that when the diarist supposedly names a sensation, say ‘S’, the conditions for this aspirant to termhood to be meaningful are not fulfilled. Why is this?

Kenny believes that the answer to this question can be given in its most basic form by posing a paradox. He writes (personal communication):

One day X has a sensation. He says ‘I’m going to call this kind of sensation “S”.’

Next day he has a sensation. (We can’t speak of ‘the sensation S’; for what is at *issue* is whether there is yet any meaning to ‘S’.)

He says ‘This is S again.’

If this is a genuine judgement, then it must be possible for it to be mistaken. If it is to be mistaken then (a) he must know what he means by ‘S’ (otherwise he won’t be mistakenly judging that it is S); and (b) he must not know what he means by ‘S’ (otherwise he won’t be mistaken).

From this it is taken to follow that the putative term ‘S’ is in reality not a term at all, but a mere meaningless noise.

The first thing to say about this argument is that, as stated, it could be taken as completely general, in which case no general terms would be meaningful. Let ‘S’ be any general term, say ‘bicycle’. When I see a man cycling down the road and say of his vehicle ‘That is a bicycle’, then if this is a genuine judgement it

²² ‘But “I impress it on myself” can only mean: this process brings it about that I remember the connection right in the future. But in the present case I have no criterion of correctness. One would like to say: whatever is going to seem right to me is right. And that only means that here we can’t talk about “right”’ (§258).

must be possible for it to be mistaken. If it is to be mistaken, then (a) I must know what I mean by ‘bicycle’ (otherwise I won’t be mistakenly judging that it is a bicycle); and (b) I must not know what I mean by ‘bicycle’ (otherwise I won’t be mistaken). Therefore ‘bicycle’ is not a genuine term.

This conclusion is, of course, absurd. Nor did Kenny intend his argument to endorse it. Fortunately, it is easy to spell out what is not made explicit in his compressed statement, and thus to guard against the possible charge that the argument is liable to over-generalisation in the way just illustrated. Kenny would say that there [45] is a difference between the predicament of a sensation term, such as ‘S’, in a private language, and the predicament of an ordinary public term such as ‘bicycle’: namely that in the latter case, but not in the former, it is possible to be mistaken without being ignorant of the meaning of the general term in question,²³ so that (b) would not apply. Only in the case of terms like ‘S’ do making a mistake and not knowing the meaning necessarily coincide. Whether the difference between the two cases is as clear as this is what I shall shortly discuss in detail: but at any rate we can now see that it would be unfair to rule out Kenny’s argument *right from the start* on the grounds that it would outlaw all general terms of whatever kind.

There is another way in which Kenny’s argument is too telescoped – whether or not its conclusion is true. It fails to distinguish between at least three different ways of being mistaken about the identity of an occurrent sensation, not all of which even *seem* to generate the *reductio ad absurdum* of the idea of a privately defined term which Kenny attempts. Besides, even in the kind of case where, *prima facie*, Kenny’s argument is valid, there is an escape route open which he has not considered.

To take the first point first. What are the different ways in which I can be mistakenly under the impression that [46] a

²³ For discussion of the various ways in which this possibility arises, see pp. 27–28 below.

SENSATIONS AS DESIGNATED OBJECTS

sensation I am having is an S-sensation? There seem to be at least the following three quite distinct types of case.

(i) My mistake may be purely verbal. For example, I recognise a sensation correctly as an instance of kind T, the kind, that is, that (in this private case) I once baptised "T", and generally mark on my calendar-diary by writing "T". But on this particular occasion I call it 'S', by a 'slip of the tongue', a more or less momentary lapse of memory as to what its proper name is. When I ask myself, or am asked, whether I am sure it is S, and stop and ponder for a moment, I realise my mistake immediately, and revise my judgement accordingly. Alternatively, I may persist in my error for the time being, though presumably not in respect of (all) future occurrences of T, since *ex hypothesi* I generally call T by its correct name.

(ii) I know the meaning of 'S' in general perfectly well, but in the particular case in question I am in some doubt. Perhaps the sensation is faint, or a borderline case, apparently ambiguous as between S and T. Eventually I make a decision, though I still remain unsure, and plump for its being S, when in fact it isn't. (There may *be* no definite answer, of course.)

(iii) I don't know what 'S' means, in general, and so I am not equipped, even in straightforward cases, to use the name 'S' correctly. I have a sensation which is properly called "T", but being under the impression, because of a [47] thoroughgoing failure of memory, that 'S' means what in fact "T" means, I call the sensation 'S'.

The crucial distinction, not mentioned by Kenny in the statement quoted above, and without which these three types of case could not be segregated, is between knowing in general what 'S' means, and knowing in a particular case that a sensation is an S-sensation. Kenny's case is set up on the assumption that my performance in one particular individual case is necessarily a definitive test of my knowledge-in-general of the meaning of 'S'. If I get it wrong, he argues, this shows *per se* that I do not know

what ‘S’ means. But this need not be so, as cases (i) and (ii) illustrate.

It may be replied that in the particular, special, hypothetical case imagined by Kenny it is in fact the case that my knowledge-in-general of the meaning of ‘S’ is necessarily tested by what I say about the (supposed) second occurrence of S. This would be so because in the situation imagined there are only two occurrences of S. The first is the occasion of the private baptism of S as ‘S’, which leaves only the second as an occasion for testing whether the lesson of that ceremony has been learned. There is no question of saying that, though on this occasion I cannot use the term ‘S’ correctly, my knowledge-in-general of the meaning of ‘S’ is sound. For there are no other occasions on which my knowledge-in-general could be evinced, tested, established, and thus given an acceptable [48] existence alongside that of my ignorance on this particular occasion.

This argument is something of a red herring, in that Wittgenstein’s specification of the diarist example (quoted above, p. 32) makes no mention of there being only two occurrences of S: on the contrary, we are clearly meant to assume that S recurs many times. And Wittgenstein nowhere else uses arguments based on a two-trial example. But it will be as well to answer the argument, both because Kenny’s specification *is* of a two-trial case, and because some of the considerations that count against the argument from such a case are also effective against the diarist example as set up by Wittgenstein.

The argument is, once again, verificationist. A non-verificationist is free to hold, simultaneously, (*a*) that I have an adequate knowledge-in-general of the meaning of ‘S’, (*b*) that I misuse ‘S’ on this occasion, and (*c*) that my knowledge-in-general has never been put to the test on any other occasion. But even without bringing verificationism into it, we can find grounds for objecting to the argument. To consider the learning of the meaning of a term in a state of arrested development after only two trials, the first of which was used for laying down the rules by which performance on future trials would be assessed, is not only

SENSATIONS AS DESIGNATED OBJECTS

artificial, but also imports the special features of the two-trial situation, which in turn create their own special difficulties, without shedding any general light on the question of the meaningfulness [49] of terms like 'S'.

For a start, there may well be future occasions when my claim to knowledge-in-general of the meaning of 'S' will be amply tested and vindicated. And even if no such future occasions arise, the same can be said of the future as can be said of the period between the first and second trials, viz. that it may be true, albeit counterfactually, that *had* there been other occasions, I should have performed perfectly well. Knowledge is dispositional, and not dependent for its existence on being actually exercised.

Further, we may charge the two-trial example with crudity. Any process of learning is liable to involve false starts, self-regulations, response to feedback and new evidence, the floating of new hypotheses in an attempt to absorb outstanding intransigent inconsistencies, and all the other devices used to guide and monitor acquisition of knowledge; and there is no reason to suppose that the establishment of a private vocabulary of sensation quality would be any exception. So to be prepared to test whether the meaning of a term has been successfully learnt after only two referents for it have been encountered is not to take the facts of human psychology seriously. Only on a wildly idealised view of learning could such an early stage of the process yield, trial by trial, a sure test of the as yet unreached upshot of the process as a whole. While learning is still going on, it is too early to impute either ignorance or knowledge. Neither successes nor mistakes can be regarded as fully fledged: [50] rather they will be embryonic, provisional. So even if, in the case of private terms, one can imagine circumstances in which error on the second of two trials entails ignorance of meaning, then this ignorance too can be only half formed. Kenny's argument will demand both half-formed knowledge of meaning and half-formed ignorance of meaning in such cases of error; and these, far from being paradoxically inconsistent though simultaneously required conditions, might more plausibly be regarded as the

same intermediate, transitional state under two different descriptions, one more forward-looking than the other. This muddying of the situation would save us from Kenny's *reductio ad absurdum* (if we needed such salvation), while doing nothing to prevent 'Here is S again', uttered on the second occasion, from being a genuine judgement, a fully-formed truth or falsehood. Approximations to knowledge are not of approximations to facts.

It seems, then, that we should not allow the special two-trial example to be regarded as an especially well favoured field of operations for Kenny's *reductio ad absurdum* of the notion of a privately defined sensation term. The moment we add a third occurrence of S, potential or actual, it becomes possible to separate knowledge-in-general of the meaning of 'S' from ability to use 'S' correctly on a particular occasion (not to mention cases of types (i) and (ii), as distinguished above). And in any case, the two-trial example is highly artificial, as just discussed. [51] Arguments which depend for their validity on appeal to such a two-trial example, then, should not detain us further.

Let us return now to the three ways of being mistaken in thinking that a sensation is S, distinguished on p. 35 above. Cases (i) and (ii) are immune to Kenny's argument. In both cases I do know in general what 'S' means, and yet I am mistaken in calling a sensation 'S'. My being mistaken in no way entails that I do not know what, in general, I mean by 'S'. In case (i) my memory or my speech mechanism fails me, and in case (ii) I am dealing with a borderline case.

It is case (iii), if any, to which Kenny's argument is germane. And this brings me to the second point canvassed above on p. 34: that there is an escape route open even where, *prima facie*, his argument seems plausible.

Suppose Kenny were right that, if I can be said to be making a genuine, though mistaken, judgement that a sensation is an S-sensation, I must know what 'S' means.²⁴ How then, in case (iii),

²⁴ If, not having a clue what a word means, I use it in an utterance correctly by accident, do I mean what I say? See the next paragraph.

can it be said that I do know this? *Ex hypothesi* I do *not* know it. Otherwise I wouldn't be mistaken.

First consider a situation where I have never known [52] what 'S' means. In the case of the private sensation vocabulary, this amounts to saying that I have never given 'S' a meaning, never used the name 'S' to baptise a type of sensation quality. If I then say 'This is S', can I be making a genuine judgement? Here at least the answer seems to be a clear-cut 'No': all I could possibly be doing would be conducting a baptism, which might indeed be infelicitous (if the sensation quality I was 'baptising' had already been given another name by me on a previous occasion), but could hardly be mistaken. (If in this predicament I said 'This is S again', then I would be talking nonsense.)

But this seems to be an exception to the general principle enunciated above (p. 32), that what holds of the private language also holds of our actual sensation language. For it is not straightforwardly true of our actual sensation terms, or indeed of any ordinary general term, that not knowing (exactly) what it means entails not being able to make a genuine judgement that some item falls under it. If I have not the faintest clue what sort of term a term is, then perhaps, like a parrot, I cannot mean anything by it in the ordinary way. But there are many terms of which we surely do know the meaning, without ever having been able to use them to identify suitable referents. I know what 'migraine' means, but I have never had a migraine, nor observed one in action on anybody else; so I would be unable, possibly, to identify a migraine of my own or someone else's if I met it. I know what 'redwood' means, [53] but I could not pick out the redwoods in an arboretum. The same goes for many terms for mechanical devices, whose functions I understand without being able to recognise the devices by their physical appearance. This is a large topic which I shall return to later.²⁵ For the present I only want to suggest that, in ordinary cases, knowing the meaning of a term is not an all or nothing affair, so that we cannot say simply

²⁵ p. 51.

‘You don’t know the meaning, so you aren’t making a genuine judgement.’ There *are* cases, such as those just mentioned, where although in some sense my knowledge of the meaning of a word may be imperfect, yet it would be strange to deny that I could use the word significantly in a genuine judgement. I know that redwoods are a kind of tree, at least; that migraines are a kind of headache; that industrial knitting machines produce knitwear. So I know the appropriate contexts in which to look for these things – I will not, for instance, try to refer to a machine by using the name of a tree. To this extent, then, even if I am hazarding a guess when I say ‘That is a redwood’, my judgement is perfectly genuine.

In the special private case, though, as I have allowed, there is no analogous line to take, and ‘This is S’ will not be a genuine judgement if I have never given ‘S’ a sense. (‘S’ derives its meaning solely and entirely from me: that is what is peculiar about terms in a private language.) [54] But this concession is quite ungrudgingly made, for the interesting case is that in which I *have* given ‘S’ a sense: am I *in this case* in a position to make a genuine, mistaken judgement, contrary to what Kenny argues?

It seems to me that the answer here is ‘Yes.’ Discussion of why this is so will occupy much of the rest of this chapter, and will arise from another version of Kenny’s argument, but as a preliminary to this I shall describe one way in which my positive answer might be justified.

The case we are imagining is one in which I once named a sensation of a particular subjective quality ‘S’, and have now forgotten which is the relevant subjective quality. Perhaps I forgot this almost immediately, and so only knew what ‘S’ meant on the occasion when I invented the term, when I used it to baptise an occurrent sensation. Even this would be enough, for it makes available the possibility of defining ‘S’ as ‘the type of sensation I named “S” at *t*’. If subsequently to *t* I forget what that sensation felt like, my (mistaken) judgement that the sensation I am now having is the sensation S is still perfectly genuine. I have not forgotten, after all, that ‘S’ is a sensation term, and I judge

SENSATIONS AS DESIGNATED OBJECTS

that the present sensation is like the one I had at t (or simply that it is S). It isn't; but it is quite clear what I am wrong about, what the content of my mistake is.

It would not be open to Wittgenstein to rule out this sort of definition *ex hypothesi*. True, it is a feature of his example that the sensation in question is not placed [55] in a reliable public context in terms of which it could be identified. But a temporal context, at least, is unavoidable. Everything happens at a time, including the private baptism ceremony in dispute. Even if we feel constrained to leave the particular time t out of it, perhaps because we cannot be guaranteed to know the value of t (or indeed to remember how many previous occasions there were on which we used the term 'S', or which was the first), we can still use the definition 'the type of sensation I once/first named "S"'. I may have forgotten which type of sensation this was, and when I first had it, but if there is indeed a type of sensation that I did baptise 'S',²⁶ then this fact alone provides a standard to which future uses of 'S' either do or do not match up. I may not be able to tell whether they meet this standard or not, but this does not alter the fact that, after all, they either do or do not meet it. A standard of correctness does not have to be subjectively recoverable.²⁷

So far, then, we have found no reason not to believe that the private linguist can make genuine judgements, mistaken or otherwise, about the identity of his private experiences, even if his memory lets him down. The paradox set up by Kenny is only properly applicable to the very kind of case which, by common consent, is a non-starter as a candidate for genuine judgementhood – the case where no baptism ceremony ever took place. But Kenny has another way of making his point, which makes use of a rather different argument, and this we must now consider. If Kenny is right, [56] then even when there has been a

²⁶ [What are the boundaries of this type? How is 'S' to be taken? (Cf. 'tove'.) As a subjective quality.]

²⁷ [In all (or any? – if I guess) of its particulars.]

SENSATIONS AS DESIGNATED OBJECTS

baptism, genuine judgements about recurrences of S are impossible.

In his 1971 article Kenny appeals, as do those who support the memory scepticism view, to Wittgenstein's §265 (in which he substitutes 'the right memory' / for 'a memory which is actually *correct*' in Anscombe's translation, on the grounds that this is what the German really means). I will say in due course why I think this appeal misplaced. First I must characterise the argument which he believes is implicit in what Wittgenstein says. It goes as follows.

When I say 'Here is S again', and check my judgement by consulting my memory of what I mean by 'S', I must call up 'the right memory', viz. the memory of what 'S' means rather than, say, the memory of what 'T' means. (Whether my memory of what 'S' means is *accurate* or not is a quite separate question: all that is at issue at present is whether the memory has, as it were, the right reference.) This in itself gives me a way of going wrong: it may seem to me that it is 'S' whose meaning I am remembering,²⁸ but why should I trust my memory in this regard any more than I trusted the initial judgement which I am now trying to check by consulting my memory? Perhaps what I am really remembering (accurately or otherwise) is the meaning of 'T'.

This difficulty alone might be thought sufficient objection to the use of private terms, from either the memory sceptic's or the verificationist's standpoint: [57] the checking procedure is simply not sufficiently independent of the fallible faculty whose judgement it supposedly confirms. But worse is to come. It is claimed that the memory I summon in order to check my initial judgement 'Here is S again' is exactly the same memory which I must make use of, even if unconsciously, to make that judgement in the first place. Thus not only am I attempting to confirm one impression by another impression drawn from the same untrustworthy source, but the two impressions turn out to be

²⁸ [What determines what it is whose meaning I am remembering?]

SENSATIONS AS DESIGNATED OBJECTS

identical, and so not independent of one another even to the minimal extent of satisfying different descriptions. Kenny writes (pp. 218–19):

We are supposing that I wish to justify my calling a private sensation ‘S’ by appealing to a mental table in which memory-samples of private objects of various kinds are listed in correlation with symbols. To make use of such a table one must call up the right memory-sample: e.g. I must make sure to call up the memory sample that belongs alongside ‘S’ and not the one which belongs to ‘T’.

But as this table exists only in the imagination, there can be no real looking up to see which sample goes with ‘S’, i.e. remembering what ‘S’ means. But this is precisely what the table was supposed to confirm. In other words, the memory of the meaning of ‘S’ is being used to confirm itself. Since the process involves our making use twice over of a single memory – the memory of which sample corresponds to ‘S’ – it is fairly compared to purchasing two copies of the same newspaper.²⁹

This passage requires some unpacking before it is clear what is being said. We begin with someone having a sensation and saying ‘Here is S again.’ We may assume that he did once christen a sensation ‘S’, whether or not the assertion with which we are concerned is true or false. The question then arises how he knows, when he says [58] ‘Here is S again’, that the sensation he is having really is S and not, say, R. How can he justify calling it ‘S’? Well, he searches his memory in an attempt to discover which type of sensation is correctly designated by the sign ‘S’. His memory of which sensations are designated by which signs is likened to a table such as a colour chart on which samples of different coloured paints are labelled with the names by which they must be ordered from the dealer. He scans his memory chart for the memory sample that belongs alongside ‘S’.

²⁹ [The newspaper may be right!]

SENSATIONS AS DESIGNATED OBJECTS

Of course the scanning will be any good to him only if he fastens on the right sample, the one that does belong alongside 'S' – and not, say the one that really belongs alongside 'T'. (And even if he gets that part of it right, his memory may still be inaccurate: but Kenny is not here concerned with this stage of the procedure.) At this stage things become more difficult, but the only coherent way to interpret Kenny's second paragraph seems to be as follows.

Since the memory table exists only in the mind of the person consulting it, it cannot yield anything but a repeat performance of the mental act that underlay the initial assertion 'Here is S again' – the assertion which is supposed to stand in need of justification. If the person in question is wrong to recognise the sensation as S in the first place, then necessarily he will come up with the same wrong result when he consults his memory table. All he is doing is [59] using the same memory twice over: in the first case, when he says 'Here is S again', he has a sensation and looks for the name that fits it (or, more likely, just 'comes up with' a name that seems right, without conscious deliberation: but for all that, a memory of what 'S' means is bound to be involved); in the second case, when challenged to justify his identification of his sensation as S, he looks for the memory sample that fits the name whose meaning he is concerned to discover, or check. But whichever side of the relation he starts from, from name or sensation quality, it is a memory of the same relation that he will be using, a memory which will be equally faithful or unfaithful in either guise. If he doesn't know what 'S' means to begin with, that's it: 'the memory of the meaning of "S" is being used to confirm itself'. He is trying to verify his memory of what 'S' means by consulting his memory of what 'S' means, and thus is no further on.

In order to discuss this argument usefully, we need to distinguish more carefully than we have done so far between the various different acts of memory which can be involved in the 'Here is S again' case. Kenny's argument essentially involves claiming that certain apparently different memories are in fact the

SENSATIONS AS DESIGNATED OBJECTS

same memory doing double duty in different guises. To assess this claim, we must be quite clear what the different acts of memory are. There are more distinctions than are made explicit at the [60] starting point of the argument. There seem to be at least three relevant things to be distinguished, which I will now describe.

(a) *Recognition of a sensation as of type S*

I have a sensation, and recognise it as S – judge that it is S, if you like, though this is an unnatural way of putting it if I do not hesitate – without conscious effort or memory consultation. I will, of course, rightly be said to be using my memory. If I was not doing so, a correct recognition would be purely coincidental – indeed no real recognition at all. But my use of memory is of the very common unconscious sort involved in most manifestations of knowledge.

What is the memory in this case a memory of? Of the fact that sensations like the one I am now having are called ‘S’? Of the meaning of ‘S’? Either description seems possible, though the first perhaps a little more natural. But they do more or less, in this case, amount to the same thing. At any rate, I can’t know, in general, that sensations like the one I am now having are called ‘S’ without knowing the meaning of ‘S’, or vice versa. Herein lie the seeds of another variety of attack on the possibility of private terms, to be considered in due course. [61]

(b) *Calling to mind what S is like*

I call to mind what S is like, without this act of memory being necessarily tied to a particular occurrence of S. Indeed I may not remember any occasion on which I had S, though I do know that I have had it at least once, and I do remember what it is like. This seems perfectly possible. Compare a more ordinary case. I may know quite well what the taste of tarragon is like, and that this knowledge is based on my past experience, without having a clue

SENSATIONS AS DESIGNATED OBJECTS

when I tasted tarragon, or whether I have done so more than once.

Of course, I may also happen to remember what the first occasion was on which I had the sensation in question – or at least remember the particular details of one occasion, if not the first, on which I had it. But this is an optional extra as far as concerns the distinctions of which I shall be making use.

(c) Remembering what 'S' means

This will be discussed below.³⁰ Armed with this tripartite distinction, we can now consider Kenny's double-duty memory argument more productively. [62] First of all, it should be clear that (a) differs from both (b) and (c). In the case of (a) I am presented with an experience which I recognise as S, and can thus identify as S in my diary. In the other cases there is not necessarily any sensation occurring at the time when I perform the act of memory specified. All that needs to occur is a memory of a sensation, of one sort or another. Even if we modify (a) so that I have to rack my brains to remember what the sensation I am now having is called, there still remains a difference: I am looking for a name to fit a sensation rather than looking for a memory of a sensation to fit a name. These procedures are certainly not identical. It is true, of course, that the same faculty is employed in both, viz. my memory. And if it is the general reliability of my memory that is in question, it may be held that matters are not much advanced by considering two of its supposed feats rather than just one. Why should a (b)-memory or a (c)-memory be held to confirm an (a)-memory? I have already given my reasons for discounting this kind of general scepticism about memory. Here I will just make the additional point that it is

³⁰ There are other permutations which we shall not need, involving all the different possible combinations of remembering that a sensation occurred on some occasion, that a quality was experienced, what the occasion was, what the quality was, what the sensation's name was, and more besides. Thankfully we do not need to consider all these permutations in turn.

not clearly irrational to believe that if we come up with the same answer to a question asked in two different ways, we have at least guarded against two kinds of possible error: (*a*) the error that can arise from thinking once rather than twice; (*b*) the error that can arise from thinking from only one point of view, from looking at a problem only from one vantage-point. [63]

When in Kenny's account the person consults his memory table, it is not entirely clear whether one can distinguish between his performing an act of (*b*)-memory and his performing an act of (*c*)-memory (though in either case the difference between what he does and an (*a*)-memory is as described in the previous paragraph). So the next question is whether (*b*)-memory is distinguishable from (*c*)-memory, and if not, what follows. Kenny says 'the memory of the meaning of "S" is being used to confirm itself. That is, when challenged to justify his claim, implicit in the assertion 'Here is S again', to know what 'S' means, the diarist consults his memory table: but, Kenny claims, this boils down to remembering what (he thinks) 'S' means, and has no status as an independent check.

Well, is there any way of doing (*c*) without doing (*b*)? Certainly to do (*b*) is to do (*c*), in this special case, for the meaning of S is entirely determined by what it is like. By the same token, to remember what 'S' means is perhaps most naturally done by remembering what S is like. But is this the only way of remembering what 'S' means? Is to do (*c*) necessarily to do (*b*)?

The quickest reply is that sketched earlier (on pp. 40–41). If we define 'S' as 'the type of sensation I once/first named "S"', then it is immediately possible to distinguish remembering what 'S' means and remembering what S is like. Remembering that 'S' means 'the type of sensation I named "S"' is quite different from remembering what S feels like. [64] Of course this move provides no practical leverage: if I misremember what S is like, I shall also misremember which sensation it was that I named 'S'. So I have no independent means of discovering what 'S' means. What then is the purpose of pointing to this distinction? Only to try to sidestep the argument 'Remembering what S is like is

identical with remembering what ‘S’ means: therefore ‘S’ is meaningless.’ But since, even if the premiss of this argument is true, I do not believe that its conclusion follows, perhaps it will be best now to confront the argument head on, as if sidestepping ploys were unavailable.

Our problem, then, is that to summon the memory of what S is like is *eo ipso* to summon the memory of what ‘S’ means, and vice versa. And indeed this problem is already present before any memory summoning takes place: even the act of memory involved in case (a) could be described in either way – as ‘remembering what S is like’, or as ‘remembering what ‘S’ means’. One act, conscious or unconscious, deliberate or involuntary, seems necessarily to achieve simultaneously two results which in normal circumstances, it is claimed, can be achieved separately, either one without the other, and which can normally be tested by different means (as I explain in a moment); also, failure can normally be diagnosed as failure of one kind or the other. In the case of ‘S’, none of these things is possible. Therefore, it is argued, we are deluding ourselves if we think that ‘S’ is a genuine term.

What *is* the difference between the case of ‘S’ and [65] the case of an ordinary term, say ‘hamster’? In the case of ‘hamster’, is it possible to distinguish forgetting the meaning of ‘hamster’ from forgetting what hamsters are like? Well, to a certain extent, yes, though the two are by no means unconnected. As already discussed briefly above, one can be said to know the meaning of many terms without being able, in practice, to recognise their referents. I may know what ‘hamster’ means and yet still misidentify a guinea-pig as a hamster. Moreover, in the event of such a misidentification it may be possible to determine whether I have forgotten the meaning of ‘hamster’ or not by a simple test: I can be asked what I think hamsters are, and if I say ‘long-eared, bounding, large furry carnivores’, what I am wrong about is the

SENSATIONS AS DESIGNATED OBJECTS

meaning of ‘hamster’.³¹ But if I give a suitable account of the meaning of ‘hamster’, then it is concluded that my mistake arises from my not knowing what hamsters look like.

This distinction is impossible in the case of ‘S’, at any rate in a situation of type (iii), because ‘S’ is a private term defined solely in terms of subjective quality. To forget what S is like is to forget what S means; to forget what S means is to forget what S is like.

[66] If I am mistaken in making the judgement ‘Here is S again’, it is impossible to distinguish between making this mistake because I have forgotten the meaning of ‘S’, and making it because I have forgotten what S feels like. They are the same thing.

Suppose all these things are true.³² They are still totally *irrelevant* to the issue of whether ‘S’ is a meaningful term. That ‘S’ is meaningless *simply does not follow*.³³ That it follows is just asserted, on the basis of a presupposition about the necessary conditions for meaningfulness which we are not bound to accept. Why would I not wish to accept it?

³¹ This need not be very detailed – certainly not so detailed as to make it paradoxical that I can know the meaning of ‘hamster’ without being able to recognise a hamster when I see one.

³² It should be clear from pp. 28–29 that I do not believe them to be true. Even if the same sort of memory must be used to perform two different kinds of check, this does not show that there are not two different things to be checked. I may on the one hand remember correctly what a sensation I had earlier felt like, but misremember what name I gave it. On the other hand I may misremember the quality of the sensation, but remember correctly what name I gave it, as the following dialogue illustrates:

A You named a sensation at *t*: what did you call it?

B (*rightly*) ‘S’.

A Did it feel like the sensation you are having now? (*B is now having R*.)

B (*wrongly*) Yes.

³³ All the immense and painstaking subtlety that goes into the defence of the private language argument, subtlety that promises a fundamental reappraisal of the nature of language as its upshot, seems in the end to be beside the point.

First, the memory act in question (under whichever description) is after all either achieved successfully or not, whether or not it is two acts in one, and whether or [67] not we are in a position to know whether it is successful.³⁴ Kenny maintains³⁵ that, as he presents it, the issue is not a verificationist one about the possibility of learning or checking the use of a term – rather it is about the possibility of a term having a meaning; and that his argument that ‘S’ cannot have a genuine meaning rests on no verificationist assumptions. And yet the grounds on which he argues that ‘S’ can have no genuine meaning seem to be precisely that the supposed checking procedure involved in defence of the claim to know the meaning of ‘S’ turns out to be defective, to be not really a checking procedure at all. Why should a (*b*)-memory confirm an (*b*)-memory? If we didn’t trust the first memory, why trust the second? How do we know that what purports to be a memory of the meaning of ‘S’ is not really a memory of the meaning of ‘T’? And so forth. But, once again, the defectiveness of a checking procedure is no indication that what we are trying to check is in some sense imaginary or empty.

This is why I said at the outset that the second main form of argument against the possibility of ODM collapses into the first. *However* we interpret Wittgenstein’s argument, the fundamental reason for rejecting terms like [68] ‘S’ turns out to be that, in one way or another, we cannot trust ourselves to remember what they mean, in the absence of public checks; and therefore the terms do not mean anything.

This is also why the final answer to Kenny must be an anti-verificationist one. A (*b*)-memory can be used to check an (*a*)-memory, up to a point. Even if a (*b*)-memory and a (*c*)-memory are indistinguishable, our capability is not impaired. In a way, it is

³⁴ It is not even obvious why an (*a*)-memory needs to be tested by a (*b*)-memory, quite apart from the question of whether such a test would be sufficiently independent. Why can we not trust our initial judgement, and leave it at that?

³⁵ Kenny (1971), p. 221.

enhanced: if we remember rightly, we remember what are normally two different things both at once, without need of or use for second thoughts. As long as we are prepared to trust our memories, we are immune to Wittgensteinian attack.

What Wittgenstein said

I shall now show why I think Kenny's argument cannot be held to be an interpretation of Wittgenstein's. In a sense this is a digression, since I have already given reasons for rejecting either way of interpreting Wittgenstein. But to the extent that the controversy over private terms stems crucially from what Wittgenstein had to say on the subject, it will perhaps be worth arguing that Wittgenstein was, after all, a verificationist, and not something more elusively sophisticated. Reluctance to disagree with Wittgenstein is often based partly on a misplaced sense of reverence for the mysterious subtlety of his writing. It is better to extract whatever clear theses can be extracted [69] from his text, and reject them boldly when they are mistaken.

Wittgenstein's §265, then, seems to me to show that he is not putting forward a version of Kenny's argument, but rather advocating the verificationist fruits of scepticism about memory. His imaginary interlocutor, or alter ego, is trying to defend the subjective justification involved in checking one's judgement 'Here is S again' – (*a*)-memory – by consulting one's memory of what S is like – (*b*)-memory – while (proto-)Wittgenstein maintains that 'justification consists in appealing to something independent'. The interlocutor replies 'But surely I can appeal from one memory to another [sc. from an (*a*)-memory to a (*b*)-memory]. For example, I don't know if I have [(*a*)-]remembered the time of departure of a train right and to check it I call to mind how a page of the timetable looked.'

This is an example of a perfectly acceptable procedure: to test one memory by consulting another one which provides the same information as the first by another route. The case is offered by the interlocutor as a parallel to the 'Here is S again' case: if the

timetable case is acceptable, so, surely, must the other be. ‘Isn’t it the same here?’ he asks. That is, isn’t this timetable case just like the case of ‘S’? ‘No,’ replies Wittgenstein. Why not? Because the second memory is testable in the timetable case, presumably by consulting the timetable (though see below for Kenny’s interpretation). But in the case of ‘S’ the second memory [70] (the memory of what S is like, arrived at by consulting the mental table) would not be testable. But why should such a memory have to be testable? Wittgenstein explains: ‘for this process [sc. of appealing “from one memory to another”] has got to produce a memory which is actually *correct*. If the mental image of the timetable could not itself be *tested* for correctness, how could it confirm the correctness of the first memory?’ Here Wittgenstein seems to be pointing quite unequivocally to the limitations on testing which afflict the private case – and *only in virtue of this kind of limitation* would he say that ‘S’ can have no meaning. There is no sign of the desiderated non-verificationist argument. All we find is scepticism about the reliability of memory.

Kenny, however, says the translation should read ‘for this process has got to call up the right memory’, and this fits in with his interpretation of the passage. If Kenny is right in his amendment, however, the passage as a whole seems to become difficult or impossible to interpret at all. (Whether this shows that Anscombe’s translation is right, or that Wittgenstein seems more confused when properly translated than when mistranslated, I do not know.)

Let me be more specific. If Kenny’s translation is right, the point is no longer that the time revealed by consulting the mental image of the timetable must be accurate. Rather, to quote from Hanfling (1975), who follows Kenny: [71]

The question is whether it is indeed a timetable, or the right timetable, that is being imagined. And Wittgenstein’s point is that this question can be answered in the case of the timetable, but not in the case of pain. Suppose I tell you I am imagining a train timetable. You ask me to describe it. ‘It has four legs,’ I

SENSATIONS AS DESIGNATED OBJECTS

reply, ‘a very long neck, and it’s eating a banana.’ Evidently I have got it wrong; I have called up the wrong image. ‘That’s not a timetable,’ you say. Again, I may be imagining a timetable, but the wrong one. For example, I tell you that it is printed in red. ‘That’s no good,’ you reply, ‘that’s the Sunday timetable.’

[...] Wittgenstein’s point is that with the image of pain there is no ‘testing for correctness’, as there is with the timetable. For what can I say if I am asked to describe that image? There is no sense in saying that I have (or have not) called up the right image. And so there is no sense in the idea of using that image to help me to identify my present feeling.

The point Hanfling makes in the first of these two paragraphs is, of course, a fair one. When one consults a table, it must be the appropriate table, and one must be looking at the right entry in it, before one can use it to corroborate a judgement. All the same, if it is the point Wittgenstein is making in the sentence of §265 which Kenny amends, the interpretation of the following sentence would be at least very unnatural. Even the amended sentence itself has to bear a heavy weight of interpretation, in that ‘the right memory’ will have to mean ‘a memory image (whether accurate or not) of how the right page of the right timetable looked, and not of how some irrelevant page of the right timetable, or of some irrelevant timetable, or perhaps not of a timetable at all, looked (even if this latter memory image is a perfectly *faithful* image of whatever it is an image of)’:³⁶ surely this

³⁶ [72] Although, for the sake of argument, I do not question it in the main text, the supposition that one may intend to remember A and succeed in remembering B is dubious. It may be, rather, that the ‘referent’ or object of a memory-search is intentionally determined, so that if I offer a memory as a memory of A, then, though of course this memory may be incorrect, it is still a memory of A. What would it mean to say that it might be a memory of B? If I am trying to remember A, how can one distinguish an inaccurate memory of A and an accurate memory of B? The obvious answer to this would be that if (as in Hanfling’s example) the details given in my memory report exactly fit B, then my memory is a memory of B. Conceivably this way of describing the situation might be confirmed by an inspection of the causal, physiological process that underlies the occurrence of the memory-act in question: it might

would have been spelled [73] out more fully and explicitly if it was what Wittgenstein meant? Besides, surely ‘tested for correctness’ in the next sentence can’t be taken to mean ‘tested to make sure it’s an image (faithfulness aside) of the relevant page of the relevant timetable?’ Even if it could mean this (perhaps by dint of another retranslation?), then Wittgenstein would most implausibly have omitted to mention that, in order for the image in question to be a trustworthy source of confirmatory evidence, not only must it be an image of the right object (Kenny’s ‘the right image’), but also a faithful image (Anscombe’s ‘a memory which is actually *correct*’). To mention the second requirement without mentioning the first is understandable, in that the first requirement is entailed by the second: to be a faithful image of x is necessarily to be an image of x . But to mention the first requirement without mentioning the second is not explicable in a parallel way: to be an image of x is not necessarily to be a faithful image of x .

Even if, *per impossibile*, we waive this difficulty, still, of course, Wittgenstein appears as a sceptic about memory: about a different aspect of memory, certainly, but about memory none the less. He will now be sceptical about our ability to call to mind an image of the relevant object, rather than about the accuracy of such an image, given that it is of the relevant object. There is no sign yet that he is an advocate of an argument for meaninglessness based upon observations about supposed double-duty memories. [74]

Finally, there is the morning paper analogy. Wittgenstein writes: ‘(As if someone were to buy several copies of the morning paper to assure himself that what it said was true.)’ Here,

turn out that there was an appropriate causal link between an occurrence of B and the memory-report I am now offering as a memory of A. But the question is, how is the object of the memory determined? Is this object intentional, or is it whatever stands in a certain causal relation to the memory? If it is intentional, then even a memory causally derived from B will still be a memory of A if that is what the memory-searcher is after; and by the same token, even if Kenny’s ‘the right memory’ is the right translation of the German of Wittgenstein, it represents an argument with no force.

certainly, there seem to be grounds for a Kenny-type argument, if we take Wittgenstein to be comparing the morning paper example with testing one's impression that one is having S by consulting one's memory of what S is like. Wittgenstein would then be implying that in this private case one would be making use twice over of the same memory. But one can arrive at this interpretation only by considering the cryptic bracketed sentence in question in isolation from the rest of the section (Wittgenstein himself says nothing explicitly about what the morning paper analogy is supposed to be analogous with). However, the sentence occurs in a context which suggests, as I have been arguing, that what is wrong with the private checking procedure is that the memory image used cannot be independently tested for accuracy – not that it is the same memory as the one being tested (which after all is not in a sense a *memory* at all, though it involves memory, but a *recognition*). And, given this context, it will not do to interpret a single enigmatic sentence in a way which demands a quite different context. It would have fitted in better with Wittgenstein's other remarks (as Kenny suggests at one point – only to reject the suggestion) if he had talked of the unreliability of checking one newspaper report by reading a report in another newspaper [75] owned by the same untrustworthy magnate – or perhaps by reading in another newspaper another version of the same report (this would be possible if the report emanated from a press agency used by more than one newspaper, or from a reporter who wrote for two newspapers).

Of course, it is a pity to deny Wittgenstein the use of the morning paper analogy in its most plausible role, as an illustration of Kenny's kind of argument about the case of 'S'. But if we start with the parenthesis about the morning paper, interpret it in this optimal way, and then work backwards through the section trying to make what he says earlier fit in, we will come up against just the same difficulties discussed above. Unfortunately, it is more economical to suppose that in the morning paper case Wittgenstein is using a less than apt analogy, and simply making

SENSATIONS AS DESIGNATED OBJECTS

the point that it's no good testing your memory by racking your brains. Sometimes one's expectations as to care and precision in the choice of analogies can be too high. Wittgenstein is just using a rough analogy, which happens to have a potential use in another context. Kenny attempts to exploit this potential, but Wittgenstein does not.

Intrasubjectivity and intersubjectivity

Before closing this chapter, I want to defend the intelligibility of saying that one person's sensations are subjectively similar to, or different from, those of another. Something [76] has already been said about the question of whether our *evidence* for the hypothesis that different people have subjectively similar experiences in similar situations is sufficiently strong for us to build this hypothesis into the structure of our sensation language. But it remains to defend the sheer intelligibility of this hypothesis: for unless it is intelligible, of course, the question of how firmly it is supported cannot arise.

The disanalogy between the intrasubjective and the intersubjective case. is sometimes believed to be sufficient to make it impossible to understand what it is to speak of subjective similarities or differences between the experiences of two different people, even if it is accepted that it makes sense to suppose that such similarities and differences occur within the internal life history of a single individual. One can appeal to memory, as I have argued, in the case of within-subject comparisons: but one cannot make this appeal in the case of between-subject comparisons. Also, it is said, my experiences are *necessarily* not yours – it is inconceivable that I should have your experiences, not merely a contingent impracticability: so it makes no sense to talk of the comparability of your experiences and mine.³⁷ And it follows from this that we could not include in a theory of the meaning of sensation terms a requirement that

³⁷ [You might as well say we can't compare blushes.]

SENSATIONS AS DESIGNATED OBJECTS

different people's sensations must be subjectively similar if they are to bear the same name. What would the cash value of such a requirement be, if I logically cannot compare your experience with [77] mine, either by having the two together, or by remembering a past instance of one and comparing the memory with a current instance or the other?

I shall argue that the impossibility of comparing your experiences and mine is not logical but contingent. It resides not in the immutable grammatical fact that your experiences are yours and mine mine, but in the conceivably alterable state of nature.³⁸ The contingency in question is universal, certainly, and this explains why it is so easily mistaken for a logical necessity; it is still a contingency.

A requirement that one occurrence of a particular type of sensation should be subjectively like another can be made intelligible in the intrasubjective case by drawing attention to the fact of 'psychological continuity', the continuity of an individual's stream of consciousness, which provides the possibility of memory checks. And it is because there is no psychological continuity in the intersubjective case that doubts are raised about making a requirement of a similar kind here too. The barrier between the two cases is that psychological continuity only occurs intrapersonally. If this apparently inevitable state of affairs could be shown to be merely contingent, then the barrier could be broken down.

We take utterly for granted the way in which the psychophysical careers of persons and the psychical careers of continuous centres of consciousness are in [78] one-to-one correspondence. Our concept of personal identity, with all its logical ramifications, is based on this. Parfit has in an article (1971) on personal identity pointed to some of the conceptual consequences of imagining perfectly conceivable cases which our present notion of personal identity would simply be inadequate to

³⁸ [When they *were* compared, they'd no longer be 'yours' and 'mine', – but this is not important.]

SENSATIONS AS DESIGNATED OBJECTS

describe. To cover these cases, the notion would have to be amended. When logical necessity has to adjust itself in the face of fresh circumstances implicit even in the world as it now is, perhaps logical necessity should not be seen as the ultimate and inevitable barrier to conceivability, a role in which it is often cast.

Imagine first that the two hemispheres of my brain are transplanted into two different bodies. There is continuity of experience between the original person and both resulting people. Personal identity and psychological continuity are forced to part company. It is no longer possible to see psychological continuity as a main criterion of personal identity, though at present it does fulfil this role: for, as this example illustrates, psychological continuity, unlike personal identity, is not logically indivisible – it can branch. Surely we don't want to say that I am both resulting people? One person survives as two. Parfit suggests we talk of 'my later selves'. Here we have survival and selfhood without the implication of continued identity which the notions currently carry.

This, then, would be a case where, given that intrasubjective [79] comparisons are possible, it makes sense to say of both the resulting people that their sensations are subjectively the same as those of the original person. They can both remember sensations from the unsplit period. So, by transitivity, the sensations of the two resulting people are subjectively the same. And yet it is impossible to compare them directly.

This is one step towards the result we want. We have given an intersubjective sense to the notion of subjective similarity of sensation quality. But this was only done by appealing to a common intrasubjective origin. So it might be said that we have not dealt with those intersubjective cases where the subjects have no prehistory in common. For these cases we need to imagine not personal fission, but personal fusion. I fuse with you, body, brain and consciousness. The supposedly logically impossible intersubjective comparison now becomes possible. If possible in this case, then conceivable in all.

SENSATIONS AS DESIGNATED OBJECTS

Does this achieve our aim? I think it does. It is a way in which sense can be given to the idea of the comparison of your experiences and mine. The possible objection is this: how do I know that postfusion experiences are comparable with either set of pre-fusion experiences? Suppose the two pre-fusion experiences of a sensation were in fact different: [80] what would the post-fusion experience of that sensation be like? This rather depends on what sort of sensation we are talking about, and on how fusion occurs. Parfit's system of fusion is that compatible elements remain; incompatible ones cancel out, or compromise. Suppose the two subjectively different pre-fusion experiences are incompatible: then the post-fusion experience will be a compromise between them – either an average of the two, or an amalgam of their compatible elements. The post-fusion memory of pre-fusion experiences will be subject to a similar distorting process. So Parfit's system of fusion does not help us.

Imagine a different system. When fusion occurs, half of each pre-fusion brain is put in a new body (or in one of the old ones, or in some combination of bodies: it doesn't matter). The two post-fusion hemispheres are connected, as they are in a normal brain. So there is only one centre of consciousness, not two as in the cases of bisected brains. Now suppose the following (all possible, some true): quality of sensation is a matter of brain structure; sensory representation is contralateral (the left side of the body being represented in the right hemisphere, and vice versa); the quality of the sensation in the unitary post-fusion centre of consciousness is determined by the structure of the hemisphere in which that sensation is represented. Here we have a case when direct comparison of pre-fusion bodily sensations will be possible: if the fused person puts both feet into hot water, and [81] the two pre-fusion experiences of heat were qualitatively different, the sensation of heat will be different for each foot.

Admittedly, this system does not cater for the comparison of experiences other than unilaterally represented bodily sensations. But suppose that these other experiences are represented bilaterally: when the post-fusion person has such an experience,

SENSATIONS AS DESIGNATED OBJECTS

then if the two pre-fusion experiences were not subjectively incompatible, the unitary post-fusion consciousness will be presented with both versions of the experience at once (though these versions may not be tagged to show which their pre-fusion origin was), and will thus be able to compare them. If they are incompatible (could they be?), the post-fusion person will find himself unfeeling in that regard, perhaps, and thus know at least that the pre-fusion experiences differed, if not in what way.

I have dealt with these examples sketchily. But I hope I have said enough to suggest that the impossibility of intersubjective comparison is not as necessary as it seemed to start with. After splitting or fusion I may not survive as the same person, exactly; but the important factor for our present problem, psychological continuity, is preserved. And it was psychological continuity which justified intrasubjective comparisons. So perhaps it will now be possible to accept the claim that intersubjective comparisons are conceivable though contingently impossible.

[82]

If what I have said in this chapter is accepted, then, there is no clear a priori objection to ODM. It makes sense, so far, to suppose that one of the necessary conditions for the correct use of sensation terms is that similar sensations should feel, not only similar to the same person at different times, but similar to different people. That this condition is fulfilled is something we cannot test directly, and yet we have good grounds for our assumption that it is.

Whether my version of ODM is the true analysis of sensation language is a separate question, which we still have to consider. The task of this chapter was only the preliminary one of defending my version of ODM against certain kinds of a priori attack. In the next chapter I try to show why an initially plausible alternative version is unsatisfactory.

2 C-SUBTLE

Having defended my version of ODM against the charge that it is, in one way or another, impossible, I must now defend it against a prominent rival version. This rival (etiolated) version has been thought to avoid the worst difficulties faced by my (full-blooded) version, and on these grounds alone to be a stronger candidate for acceptance. But in addition to this it is claimed on its behalf that it presents a more realistic³⁹ picture of our sensation language, and should on that account be preferred to the over-demanding account which I have offered. The theory in question is christened ‘C-subtle’ by Pears (1971, chapter 8) in a discussion of Wittgenstein’s view of sensation language. I shall now examine this theory, adopting Pears’s name for it as usefully brief, if not exceptionally mnemonic in the present context.⁴⁰ [84]

I begin by quoting Pears’s initial characterisation of C-subtle as it applies to pain, and of its most important consequence for those who find my version of ODM unacceptable:

³⁹ [In the ordinary sense of the word.]

⁴⁰ C-subtle is so called by contrast with ‘C-crude’, an even stronger version of ODM than mine (‘C’ means roughly the same as ‘ODM’) which makes things unnecessarily hard for itself by neglecting the wide range of publicly observable phenomena through whose regular association with sensations we are enabled to teach the use of sensation terms to children, and indeed to understand sensation language ourselves. Even without this deficiency, though, C-crude would be opposed to C-subtle: it would still require, as C-subtle does not, that the quality of any given sensation be intersubjectively similar. (If it not only recognised the pedagogic role of public phenomena, but also allowed that such phenomena featured in the meanings of sensation terms, it would then be the same as my version of ODM.)

the word ‘pain’ simply means ‘sensation of a type that has such and such teaching links’.⁴¹ So if, unknown to us, and perhaps unknowably to us, the same teaching links pointed to different kinds of sensation in different people, those differences would not be picked up in the meaning of the word ‘pain’. (p. 152)

So, expressed in general terms, C-subtle is the theory that a sensation term ‘*x*’ means ‘sensation of a type that has such and such teaching links’.

The first thing to notice about this formulation of the theory is that it is ambiguous. On one interpretation, the theory claims that what makes a sensation the sensation *x* is simply and solely that it occurs in appropriate public company. Given only that it is indeed a sensation, its intrinsic quality is entirely irrelevant, and may indeed be different on every separate occasion when it occurs, so long as the links are right. Thus its ‘type’ has nothing to do with its felt quality. [85]

On the other interpretation, the ‘type’ is not exclusively a matter of links: it is a subjective type too, and for any given person, *x* must have the same intrinsic quality whenever it occurs. But *which* quality it has *is* determined by the links, in the following way: for any given person, *x* has whichever subjective type of sensation quality⁴² happens, for that person, to occur in regular

⁴¹ ‘Teaching Links’ is Pears’s name for the public phenomena mentioned in the preceding note. He refers to them in this way because he is particularly concerned, in the context in question, with the problem of how exactly sensation language is taught. The links are, of course, not used exclusively in teaching situations: but I hope it will not be seriously misleading if I adopt Pears’s phrase in what follows. Like ‘C-subtle’, it has the virtue of brevity. What the teaching links of sensations are is discussed further in chapter 3: to mention just a few possibilities by way of illustration, a link may be the natural expression of a sensation, its usual cause, its temporal pattern, its non-causal harbinger, or its non-expressive contemporary or sequel.

⁴² There is a difficulty about expressing oneself neutrally in this area. The word ‘sensation’ may, depending on one’s viewpoint, be taken to mean either (a) something defined purely in subjective terms; (b) something [will this do?] defined purely in objective terms; (c) something defined by a combination of both. Depending on which view is true, one will express oneself differently.

company with the links which are definitive of x ; the quality of x is just the quality the sensation so picked out happens to have. It may be one quality for me and another quality for you (if it is allowed that this means anything): but as long as it is the same quality all the I time for any individual, the demands of C-subtle are met.

The first interpretation yields a highly implausible theory, according to which the only reliable way of telling what kind of sensation is occurring is to observe its causes or effects. Such a test would have to be applied on every single occasion on which a sensation occurred, throughout life, not only while the meaning of sensation terms was still being taught. If one used the intrinsic quality of a sensation as a guide to what sort of sensation it was, one would be taking the risk of following a mere rule of thumb whose applicability depended on a contingent connection between sensation quality and links, a connection that had nothing to do with meaning.

This is such a caricature of the way we all operate, and know that we operate, with sensation terms that it is not perhaps worth spending much time on it. But it must be said that it is not a logically impossible view. It makes sense, at least, to suppose that

For example, if one wishes to refer exclusively to subjective quality, then if (a) is true one need only say 'sensation', if (b) is true one may be tongue-tied, and if (c) is true one needs to say 'subjective quality of a sensation' or 'sensation quality'. The question 'What makes a sensation the sensation x ?' should, ideally, be debated in theory-neutral terminology. The purpose of this note is to explain, lest I should be charged with prejudging the issue, why this is impossible. I have tried to speak explicitly of sensation quality wherever I mean to refer to the intrinsic nature of a private event as opposed to any public manifestations that event may have. (Notice that I had to use 'private event' instead of 'sensation': even 'the intrinsic nature of a sensation' is ambiguous, meaning either 'the intrinsic nature of a particular (subjective) type of private event' or 'the intrinsic nature – perhaps different from instance to instance – of whatever private event occurs in company with such-and-such teaching links'.) Unless one starts talking in terms of 'sensation₁' and 'sensation₂' or somesuch, which would be cumbersome, some ambiguity [or, rather, lack of theory-neutrality] seems inevitable.

we might talk of sensations according to this model. But it would be, to put it at its lowest, very odd to do so, in view of the fact that sensation quality *does*, experience shows, behave in a predictable way. It does keep constant company with the same links throughout a person's life, by and large. It does not vary in some totally random way from instance to instance of the same public context. Given that this is so, it would be natural to expect that we would make use of the correlation between links and sensation quality in our sensation language.

Why should we not do so? Perhaps it may be said that [87] it is rash to rely on a contingent connection of this kind: better to construct our language so that it is immune to any possible mutation in the regularity of nature. But why should we wish to build such immunity into the logic of our sensation language? Nature is not in fact so mutable as to require it. And even if regularity of the correlation between sensation quality and teaching links did collapse, we could not effectively protect ourselves in the way suggested: a breakdown of natural law as severe as this would have consequences too far-reaching to leave sensation language untouched, oasis-like amid the fundamental restructuring of language that would surely occur.

One other point should be made about this extreme interpretation of C-subtle before we leave it. If the arguments against the possibility of (partly) privately defined terms were accepted, it is this interpretation, if any,⁴³ which would have to be

⁴³ I say 'if any' because the whole notion of the subjective 'quality' of a sensation may be rejected. In which case no form of C will be acceptable, even one which makes no stipulation about what the subjective quality of a sensation should be. Even the extreme version of C-subtle which is here under discussion is a version of ODM (albeit a maximally attenuated one), in that it still pictures sensation terms as referring to private events. But by abandoning all requirements, intrapersonal and interpersonal, about type or constancy of sensation quality, it does manage to avoid what the private language argument sees as the worst excesses of ODM, and so might be a candidate for acceptance by those supporters of the private language argument who do not follow the argument through to its final logical conclusion, but allow that private events exist, so long as they do not lay claim to any substantial

adopted. The other [88] interpretation includes a requirement that, for a given person, a sensation of type x should have the same intrinsic quality on different occasions. But just this sort of requirement would be ruled out as senseless by adherents of the private language argument. So to adopt the milder version of C-subtle is already to accept an important part of what the private language argument rejects. (It is not to accept all of it, as my version of ODM does, for no *interpersonal* requirement about comparability of intrinsic sensation quality is made.) To this extent, C-subtle in its milder form cannot be claimed to be a significant improvement on my version of ODM, if the point of trying to improve on this is to avoid vulnerability to the private language argument. But it may be thought that it is the interpersonal requirement rather than the intrapersonal requirement which is the real weakness of my version of ODM.⁴⁴ In which case the moderate interpretation of C-subtle is, *prima facie*, a promising alternative, and should be examined.⁴⁵ This I shall now do. [90]

linguistic role. (Even in C-subtle they still play a role: x must still, after all, be a *sensation* at least.)

⁴⁴ Though I show above, pp. 42–46, why I reject this differentiation between the boldness of the two requirements.

⁴⁵ There is another issue, quite separate from the issue of whether C-subtle is true: viz., whether Wittgenstein held it. Geach (1957, pp. 3–4) holds, in effect, that Wittgenstein did subscribe to C-subtle: ‘what Wittgenstein wanted to deny was not the private *reference* of psychological expressions – e.g. that “pain” stands for a kind of experience that may be quite ‘private’ – but the possibility of giving them a private *sense* – e.g. of giving sense to the word [89] “pain” by just attending to one’s own pain-experiences, a performance that would be private and uncheckable.’ Locke (1968) takes a similar line. Pears (1971, chapter 8) leaves the question open.

It seems clear from §§273–5 that Wittgenstein rejected C-subtle. As he says in §274, ‘saying that the word “red” “refers to” instead of “means” something private does not help us in the least to grasp its function’: though there is nothing in his subsequent remarks which could count as an *argument* for his rejection of this way of talking.

Arguments against C-subtle

There are three main arguments against C-subtle⁴⁶ on which I shall concentrate. The first concerns a paradox which arises from C-subtle's comparative indifference to the intrinsic quality of a sensation. The second shows that C-subtle is in fact logically committed to the very sort of interpersonal subjective requirement which it is one of its main purposes to avoid. The third is based on an examination of the problems attached to applying C-subtle to the generic term 'sensation' analogously to its application to the specific term 'x'. I take the first argument first.

According to C-subtle, it should be possible to teach a child (or for that matter an adult) the meaning of the name of a sensation which he has never experienced. All he will discover when (if ever) he does come to experience the sensation is what it is like (for him). Why do I say this?

The term 'x' is defined as 'sensation of a type which has such and such teaching links'. We may assume for the purposes of this argument that the person who is to be taught the meaning of 'x' understands the word 'sensation' [91] on the basis of his knowledge of other sensation words. He can also, of course, understand about the teaching links, since they are public and can thus be described or demonstrated to the pupil – though he will in all probability already be acquainted with many of them from his observation of other people. So he can satisfy all the conditions of understanding the meaning of 'x', even though he has never had x.

That this situation can arise creates the possibility of what seems, at any rate, to be a paradox. When the person who has learned the meaning of 'x' without having had x eventually does

⁴⁶ By 'C-subtle' I shall henceforth mean the milder version, unless otherwise specified.

have x , he may not know that x is what he is having. More starkly:

- (i) He knows the meaning of ‘ x ’.
- (ii) He has x .
- (iii) He does not know that he is having x – does not recognise the sensation he is having *as* x .

This situation could arise if the links which are definitive of x are not of a kind to be easily noticed straight off by someone who is having x for the first time. So although in theory he is in a position to recognise the first occurrence of x by noticing the relevant features of the context in which it occurs (even this is a highly artificial recognition strategy!), in practice this may prove impossible. For example, if the link for [sensations of] cramp is that a certain kind of muscular spasm is occurring, the child who feels cramp for the first time is unlikely to recognise this kind of spasm for what it is. [92]

It would also be possible for a new sensation to occur without its links [or without its links being obvious], unless C-subtle rules out the possibility that the correlation between a sensation of particular quality and a particular (set of) teaching link(s) has exceptions. Experience shows that this possibility can be realised: the pairing in virtue of which (according to C-subtle) a sensation of a particular quality is the sensation x for a particular person at best holds most of the time. There are cases where, for one reason or another, it breaks down.⁴⁷ So we might modify (or perhaps it is only ‘unpack’ or ‘spell out’) C-subtle slightly, and define ‘ x ’ as ‘sensation of a type which *normally*⁴⁸ has such and

⁴⁷ [Examples needed? Spontaneous neural discharge? Only counts if links are not necessary *in every case*.]

⁴⁸ How many exceptions are allowable before the basis for defining a sensation on C-subtle’s model breaks down is an issue which I do not pursue, though it may provide another way of showing that C-subtle is unsatisfactory. [I doubt it: unless the demonstration argued that it would be impossible to have a C-subtle-type term for a sensation-quality that kept irregular company.]

such teaching links'. Then we can suppose that our child, who the time. There has learnt the meaning of 'x' without ever having had x, subsequently experiences x for the first time *without* the teaching links which normally accompany it. In this case he will not recognise the sensation as x, for he will have no means of doing so – not just means which he cannot recognise, as in the previous example.

Is this really a paradox? How implausible is it to [93] hold that a person may know the meaning of a sensation term without being able to recognise the sensation it refers to when it (first) occurs? Does this depend on what value we give to the variable 'x'? If we consider 'headache' or 'cramp', is the situation less paradoxical than if we consider 'pain' or 'red'? And if not, why the difference?

This issue is in fact quite general, and not confined to sensation language. I have already said something about it in chapter 1 (pp. 39–40), and shall now say a little more. The general question is this: how much direct [i.e. first-hand] personal experience of the world do we need in order to understand a given term or type of term?⁴⁹ How far can experience of one kind compensate for deficiency in experience of another kind, enabling us to understand new terms by analogy, by combination of elements of terms we already know, by extrapolation, imagination, guesswork or any other device?

John Locke's simple and complex ideas come to mind as tools to use in answering this sort of question. Simple ideas of sense, it seems, cannot be known except by [94] acquaintance.⁵⁰ So if

But any kind of public term for this would be a problem, surely?] Another relevant issue which I sidestep here is the whole question of criteria, of the logical relation between teaching links and meaning. Something is said about the simpler aspects of this in chapter 3.

⁴⁹ The question how much experience is required for an understanding of language in general is a different and far more difficult one which I do not consider.

⁵⁰ 'When the understanding is once stored with [...] simple ideas, it has the power to repeat, compare, and unite them, even to an almost infinite variety,

there are terms in our language whose meaning is wholly or partly determined by a requirement that they refer to particular simple ideas, there will be no obvious substitute for experience as a means for understanding the meaning of these terms, except in so far as there is some specifiable analogy between a simple idea that has been experienced and one that has not.

But most terms in our language are complex: in Locke's terminology, they stand for complex ideas. Here the role of experience is not so straightforward. A complex idea may be made up of elements some or all of which we are acquainted with. If we are not acquainted with what is denoted by the relevant complex term, can we nevertheless claim to understand this term on the strength of our understanding of, as it were, its simple constituent parts?

The answer here cannot be an all or nothing one. There are degrees of understanding. Moreover, it will vary from case to case, being affected by the nature of the particular term in question, and by how comprehensive or [95] exhaustive the knowledge is out of which the understanding of the new term is to be constructed. This is where the remarks made in chapter 1 come in. The person who knows that redwoods are trees (and perhaps also that they are large, that they are reddish in colour, and one or two other details), but cannot recognise redwoods, may yet be said to know what 'redwood' means. His experience of the world is sufficient for us to grant him such understanding, by proxy. Again, in the case of a piece of industrial machinery defined by reference to its function, it is not necessary to be able to recognise the machine when we see it in order to understand the meaning of the machine's name. This is partly because the definition of the name is in terms of function and not in terms of appearance: but also – and this is why we understand the

and so can make at pleasure new complex ideas. But it is not in the power of the most exalted wit, or enlarged understanding, by any quickness or variety of thought, to *invent* or *frame* one new simple idea in the mind.' Locke (1690) 2. 2. 2.

functional definition – because we know by experience what the components of the functional definition mean.⁵¹ For example, there [96] is a machine that makes wood-screws. If there were a special name for this machine (perhaps there is), we could understand its meaning for just the same reason that we can understand the meaning of ‘machine that makes wood-screws’: we know what machines and screws are, and we know what it is for a machine to make something. But we might well fail to recognise an idle screw-making machine if we came upon it cold – without contextual clues as to what it was.

So when we wish to know whether we can be said to understand a new term without direct experience of what it denotes, the question to ask is whether the descriptions that can be offered in explanation of the term’s meaning ‘reduce’ that term to a conglomerate all, or a sufficient proportion, of whose components we either understand on the basis of our own experience, or can plausibly be said to understand through the availability of a sufficiently close analogy with something of which we do have experience.

Both of these routes to understanding come into play in the case of sensation terms, but in different proportions for different terms. And because of this difference of proportion, different answers will be produced to the question ‘Can we understand the

⁵¹ There is also what Putnam (1973) calls ‘division of labour’ in the understanding of a term. His article is very clear and plausible on this whole topic, and the points he makes should certainly be added to what I say in order to fill out the picture of how knowledge of meaning without capacity to recognise is possible. Here is what he says about ‘gold’: ‘everyone to whom gold is important for any reason has to acquire the word “gold”; but he does not have to acquire the method of recognising whether something is or is not gold. He can rely on a special subclass of speakers. The features that are generally thought to be present in connection with a general name – necessary and sufficient conditions for membership in the extension, ways of recognising whether something is in the extension, etc. – are all present in the linguistic community considered as a collective body; but that collective body divides the “labour” of knowing and employing these various parts of the “meaning” of “gold”.’ Putnam (1973), p. 705.

meaning of “*x*” without having experienced *x*?, depending on what value of ‘*x*’ is under consideration. For, first, it is more important in the case of some sensations than it is in the case of others that they should have a particular definite felt quality, and for this reason there will tend to be a [97] proportionately larger or smaller gap in our understanding of the meaning of a sensation term whose referent we have not experienced. And, second, some sensation qualities that we have not experienced will be more similar to those we have experienced than will others. We might say, then, generally speaking, that the degree of understanding of a new sensation term that we can achieve is a joint function of (a) the predominance in the meaning of that term of the requirement that the sensation have a certain quality, and (b) the novelty, for us, of that quality.

For example, if ‘cramp’ is defined as ‘pain caused by a certain kind of muscular spasm’, we can fully understand the meaning of ‘cramp’ without ever having had cramp. We know what pain is, including what it feels like (this too, I hold, is part of the meaning of ‘pain’), and we can be brought to understand about the relevant kind of muscular spasm in terms which are already in our repertoire: so our understanding of all the components of the definition of ‘cramp’ stands in for our lack of direct experience of cramp itself. Even if we expand the definition of ‘cramp’ a little so that it incorporates a reference to a special felt quality (we might for instance define ‘cramp’ as ‘pain of that peculiar subjective quality caused by a certain kind of muscular spasm’), still we may claim to understand the meaning of ‘cramp’ without having had cramp: we know what pains of other varieties feel like, and this knowledge stands in, perhaps [98] to a sufficient extent, for our lack of experience of what cramp itself feels like. All we are missing is knowledge of which particular, peculiar type of pain quality belongs to cramp, and the similarity of this pain quality to experiences we *have* had is sufficient to allow us analogical understanding of ‘cramp’. Similar things might be said about the case of ‘headache’.

To return to the possible paradox set up above: to the extent that we grant someone understanding of the meaning of ‘*x*’ without his having ever experienced *x*, to that extent there is no paradox. Rather it is true of many (perhaps most) things, certain sensations among them, that we may not recognise them on first encounter, even though we are able to talk about them perfectly coherently (if not perhaps sympathetically, in the case of the sensations) without encountering them at all.⁵²

But suppose, now, that someone has never felt pain *of any kind*. Will his knowledge merely of what sensations in general are like be enough to enable him to learn what ‘pain’ means? Hardly. In fact, if he has no experience of pain, we may be reluctant to concede even that his general knowledge of the nature of sensations is adequate. Would we say of the man with monochrome vision that he had a good general knowledge of the nature of visual experience? However this may be, pain is not an experience whose quality can be communicated by analogy to those who have not felt it, and if we hold, as I do, that an important part, at least, of knowing the meaning of ‘pain’ is knowing what pain feels like, we will not be prepared wholeheartedly to ascribe knowledge of the meaning of ‘pain’ to someone who has never felt pain.⁵³ Of course, this is not an all or

⁵² Equally, it is precisely these things that we are more likely to recognise on first encounter, for the same reason that we can understand their names before we encounter them, viz., because they are similar to or constructed out of [or analysable in terms of] things we *have* experienced.

⁵³ [Begg question in favour of my version of ODM (ODMH), doesn’t it? C-subtle is still *possible*, though indeed counterintuitive. Or is it? Perhaps I should say: ‘How can we attribute knowledge of meaning to someone who cannot use the word in question in one of its most centrally appropriate contexts?’ But wouldn’t the same apply to ‘redwood’ if I can’t identify redwoods? The difference is that sensations are crucially felt and identifiable by people in a way that redwoods aren’t. We don’t ‘divide the labour’ of understanding the meaning of sensation terms.

Wouldn’t the same objection apply to ODMH? After all, even though experience alone can tell us *what* quality is required, is this really to do with the *meaning*? Not in the sense of ‘meaning’ whereby it’s equivalent to ‘verbal

nothing affair. It can quite intelligibly be said that I partly know the meaning of ‘pain’ even if I have never experienced pain myself: otherwise how should I be able successfully to pick out which people are in pain, to the extent that I could if I had never felt pain?⁵⁴ I could learn about the public criteria of pain as well as anyone.) For the same reason, it would be more paradoxical in the case of ‘pain’ than in the case of ‘cramp’ or ‘headache’ to claim full understanding of meaning without experience of any sample referents, and so without the guaranteed capacity to recognise such a referent on first encounter; and a theory which accepts such a claim in the case of ‘pain’ is at best highly implausible. But the paradox can be avoided by disqualifying the first move towards it, the move of allowing that the person who has never felt pain can nevertheless understand the meaning of ‘pain’ as well as anyone else.

We may be inclined to disqualify, likewise, allowing that someone who has never seen red can yet learn the [100] meaning of ‘red’. Suppose we show such a person some poppy seeds. ‘Red’, we say, ‘is the colour of the flowers that will grow from these seeds.’ Or else we might say, less practically but perhaps⁵⁵ more securely, that red is the colour of light whose wavelength is in the vicinity of 620 millimicrons. Neither of these methods of teaching will enable our pupil to identify his first experience of red, if it occurs without his knowing that the links for red are present. He may see some red flowers, but not know that they grew from the seeds he was shown earlier. He may see a patch of light in a physics laboratory, but not be aware that its wavelength is 620 millimicrons. If we believe that the phenomenal quality of an experience of red is an essential part – perhaps the only

definition? – but if meaning has to do with use, indefinables may enter into it. This is also so if Kripke is right – acquaintance with samples lies at the root of meaning.]

⁵⁴ [Would *this* ability be impaired *at all*?]

⁵⁵ Is it contingent or necessary that light of wavelength 620 milli-microns is red? Even if it is contingent, it may be a safer bet than that certain ungerminated seeds will produce red flowers. See further pp. 65–66 below.

essential part – of its identity as an experience of red, then we shall not allow the possibility of fully understanding the meaning of ‘red’ without having seen red.⁵⁶

If the case of ‘red’ is less than wholly convincing, the case of ‘colour’ should fill in whatever is missing. Could a person with monochromatic vision understand the meaning of ‘colour’? This is like the case of pain: no basis for analogy appears to exist among other sensations. [101] C-subtle might offer as a definition of ‘sensation of colour’ something like ‘sensation caused by differential response of the cones to a specific wavelength of light’. But this surely isn’t a sufficient characterisation of the nature of sensations of colour. In the case of a general sensation term like ‘(sensations of) colour’ or ‘pain’, which excludes the possibility of analogical understanding precisely because it covers all sensations that fall within a certain category, we cannot be said to know the term’s meaning fully in advance of having experienced sensations that fall under it.⁵⁷ And in the case of more specific [102] sensation terms such as ‘cramp’, ‘headache’

⁵⁶ [Would we say something different for ‘red’ and ‘sensation of red’? It depends whether what makes something red is that it looks red, or not. If so, ‘red’ will not be understood without experience. If not, it could be, perhaps (what, though, in that case would make red things red?).]

⁵⁷ Cf. two passages from Kripke (1972): ‘we identify heat and are able to sense it by the fact that it produces in us a sensation of heat. It might here be so important to the concept that its reference is fixed in this way, that if someone else detects heat by some sort of instrument, but is unable to feel it, we might want to say, if we like, that the concept of heat is not the same even though the referent is the same’ (p. 325) If this is true of ‘heat’, which is not the name of a sensation, how much more will it be true of a sensation of heat that someone who cannot feel it cannot understand the meaning of ‘sensation of heat’ in the same way as others? [Still more so for ‘pain’.] ‘[T]he way the reference is fixed seems overwhelmingly important to us in the case of sensed phenomena: a blind man who uses the term “light”, even though he uses it as a rigid designator for the very same phenomenon as we, seems to us to have lost a great deal, perhaps enough for us to declare that he has a different concept. [...] The fact that we identify light in a certain way seems to us to be *crucial*, even though it is not necessary’ (p. 331). Once again, a fortiori the blind man will have an even more divergent concept of ‘sensation of light’.

and perhaps ‘red’, we must at least have experienced something analogous, if not the sensation in question itself.⁵⁸

To sum up: if C-subtle is true, we can understand any sensation term, whether or not we have experienced a sample of the kind of sensation to which it refers. Against this, our intuitions require at least experience of analogous sensations before we grant the possibility of such understanding. If these intuitions are right, C-subtle is inadequate, since even in the case of general sensation terms it does not require experience as a prerequisite of understanding. ‘Pain’ will, according to C-subtle, be defined in just the same way as other sensation words, as ‘sensation of a type which has such and such teaching links’, so that acquaintance with the intrinsic quality of pain is not a necessary condition of understanding the meaning of ‘pain’. Ultimately, however, understanding is grounded in experience: less so, perhaps, in specific instances which can be fitted into general contexts that are already understood, but by no means all sensation terms, as we have seen, can be excepted under this rubric.

If C-subtle must be rejected in some cases, and so becomes unavailable as a general model of sensation language, it may be better to abandon it entirely, for the following reason. To adopt C-subtle for any term, even tentatively, is to assume that the meaning of that term is innocent of a [qualitatively specific] requirement about subjective quality until proved guilty. That is, it is to begin from a prejudiced position. **[103]** There may indeed be some terms whose correct analysis will be of a form that is compatible with C-subtle: but whether this is so should be left

⁵⁸ [This argument isn’t really directed straight at C-subtle, which might *agree* that experience is necessary. *What quality* of experience would be immaterial, of course, since it is the links that matter, but there’s no reason why C-subtle shouldn’t hold, like any other theory, that experience must precede full understanding – or is there?]

Well, there is the point that ODMH requires a particular definite quality – so experience is obviously required; but C-subtle requires only *some definite quality or other*, and experience is conceivably dispensable.]

open at the start of any particular enquiry. A better model will have a slot for [a specific (intrinsically)] subjective quality⁵⁹ and a slot for teaching links, either of which may turn out to remain unfilled:⁶⁰ but neither slot should be barricaded a priori.

Also, unless C-subtle is universally applicable [i.e. to all sensation terms], there is no longer any general theoretical justification for importing it as a substitute for my version of ODM. If the motivation for introducing C-subtle was to provide an across the board replacement for a theory which was supposedly inadequate or incoherent, then the need to reject C-subtle in some cases loses it its prime appeal: it can no longer be an all-inclusive theory of sensation language, but only (if anything) an analysis of some sensation terms.

Before leaving this argument against C-subtle and proceeding to the next one, there is a related oddity to which attention may be drawn. As explained above, if C-subtle is correct, a person does not need to know what x feels like in order to understand the meaning of ' x '. At least, this applies if he has not yet experienced x . But does it continue to apply once he has experienced x ? It seems not, for the following reason. C-subtle requires that, for a given person, x must have the same intrinsic quality whenever it occurs (though which quality this is is decided by reference to **[104]** the relevant teaching links, as explained). This requirement reflects the fact that we recognise our sensations by the particular way they feel (however irrelevant this may be to their identity as x s), so that if they stopped feeling the same on subsequent occasions, we'd be in trouble. We tell, in most cases, that we have sensation x because we know what x s feel like (for us at least), and the sensation we are having feels like that.

Our recognition of our sensations, then, depends on our knowing what they feel like. And of course there would be no

⁵⁹ [i.e. not just 'whatever sensation happens ...?']

⁶⁰ [Not links, in a public language? Yes – they may be reference-fixers, not part of meaning.]

point in C-subtle's requirement that the intrinsic quality of x should be intrasubjectively constant unless we did recognise sensations in this way. If we could not tell whether the quality of a sensation was the same or different on different occasions, the requirement would be unenforceable, and so useless and irrelevant. But it follows that, once I have experienced x , I cannot know the meaning of 'x' without knowing what x is like. C-subtle requires that x always has, for me, the same quality. Before I experience x , what this quality should be is open. After I have experienced x , what it should be is fixed. For a sensation to be the sensation x , thenceforward, it must have *that* quality.⁶¹

Would it not be extremely odd if one of the of knowing the meaning of a word applied after, but not before, one had encountered an appropriate referent for the word? It cannot be claimed that this oddity shows that [105] C-subtle is straightforwardly impossible: for there is no a priori reason why the conditions of knowing the meaning of a word should remain constant for a given person.⁶² But it does make it look as if x means one thing before I have it, and another afterwards.⁶³

Considerations of this kind also support a possible argument against the absence in C-subtle of any interpersonal requirement that sensations of a certain kind be subjectively similar, at least if we accept certain arguments of Strawson's in *Individuals* (1959).

The relevant points Strawson claims to establish are these: it is logically necessary that the (logically primitive) concept of a person should be the same whether it is applied to oneself or others. It is learnt *as* such a dual-purpose concept – not learnt in one context and then extended to suit the other. Unless it applied equally and in the same sense both to oneself and to others, it would not be the concept it is. From this it follows that 'P-

⁶¹ [Cf. 'my favourite cocktail', if this is not allowed to vary over time.]

⁶² [Yes there is: see next sentence.]

⁶³ [Is this really a good argument? Cf. 'the winner of the 1977 election' – ah, but here even afterwards it's not part of the meaning ...]

predicates' (predicates peculiarly applicable to persons, including sensation predicates) essentially have first as well as second and third person uses. To learn their use is to learn both aspects of their use. It is a necessary condition of one's ascribing experiences to oneself as one does that one should also ascribe them to others, and *in the same sense* (cf. figure 1, p. 18).

But, as Strawson notes, there are some P-predicates, sensation predicates among them, which I ascribe *to myself* on the basis on non-behavioural evidence,⁶⁴ but *to others* on the basis of behavioural evidence. In which case, if Strawson's conclusions as just summarised are correct, we must suppose that others do *precisely the same*. If we look carefully at just what this requirement of conceptual parallelism involves, we shall find that it is fatal to C-subtle, for it commits us to that very intersubjective rule of phenomenal quality which C-subtle omits. How is this so?

In one's own case it is required by C-subtle that occurrences of sensation x should all be subjectively similar. So far, so good: C-subtle imposes this requirement in the case of others too. But there is a further requirement that one makes in one's own case, which C-subtle does not want held in common. This is the requirement that all occurrences of sensation x should feel like *this*: not just that they should be similar in some unspecified way, but that they should be similar in virtue of sharing some particular subjective quality. As I have argued above, it seems to become part of the *meaning* of ' x ' in one's own case that just this subjective experience is present. Now, since I must, once again, ascribe P-predicates to others⁶⁵ in the same sense as that in which I ascribe them to myself, this extra requirement must apply to the case of others as well. If it does not apply, then there is no guarantee that the meaning of ' x ' for others will be the same as it is for me. For although they will need to have the same subjective experience on all occasions of [108] the occurrence of x , it need not be subjectively similar to the experience I always have when I

⁶⁴ [Or on no evidence? Cf. his p. 107.]

⁶⁵ [Or 'hold that others ascribe P-predicates to themselves'.]

have x . One of the necessary conditions for the applicability of ‘ x ’ will be different in the two cases, and Strawson’s requirement will not be met.

It will be objected that the interpersonal requirement I have just been discussing is deliberately excluded from C-subtle, and that it should not be sneaked in again. There are, after all, as many different ways of feeling the same as there are different sensation qualities, and the requirement that C-subtle does make, that, for any given person, x should feel the same on all occasions, does not entail that it should feel the same for everybody. I have simply confused sameness-for-a-given-person with sameness-for-all. The requirement of sameness-for-a-given-person is the same for myself and others, so C-subtle does pass the Strawsonian test.

Why do I not accept this objection? Because in my own case, as I have argued, once I have learnt by means of the appropriate teaching links which subjective kinds of sensation to pick out as the proper bearers of which names, then the way a sensation feels becomes a criterion of its identity. Once *this* type of sensation quality has been settled on as properly called ‘ x ’, it is no longer true that there is any number of ways for all occurrences of x to be similar: there is only one way, viz. to be similar to *this*. That a sensation feels like *this*, then, [109] is now a necessary condition for the applicability of ‘ x ’ in my own case – and is part of the meaning of ‘ x ’ for me. And unless the same condition applies in everybody’s case, Strawson’s condition is not met.

Of course, it may be maintained that it is not part of the meaning of ‘ x ’ just what the sensation feels like: only that it always feels the same for a given person. Perhaps it will be said that what it feels like is a matter of reference rather than of meaning. But this is arbitrarily to exclude one criterion for the applicability of a term from playing a role in that term’s meaning.⁶⁶ [110] There are indeed expressions which refer to

⁶⁶ This is a very seductive move to make. Pears writes (p. 154): ‘Someone who wanted to defend C-subtle would claim that, though the connection

particulars under descriptions which do not allude to the particulars' particular properties. To be queen of England one does not have to be the same person as some particular queen of England. But if C-subtle is accepted, then a sensation, to be the sensation x , *does* have to have a particular subjective quality, not merely be the associate of certain teaching links.

Again the objection will be pressed. Just as two people can be the same qua queen of England, but different qua individual persons, in the same way, surely, two types of sensation can be the same qua associates of such and such teaching links, but

between reference and meaning is close, it is not so close that, if the reference is private, the meaning must be private too. He would point out that everything depends on the way in which the private reference is made, and that in his theory it is made in a way which ensures that it does not contribute to the meaning of the phrase "sensation of a type which has such and such teaching links", just as the actual reference of the phrase "a person's favourite cocktail" does not contribute to its meaning.' But, as I argue in the text, this works only for the extreme version of C-subtle. My favourite cocktail can change from one year to the next: but if all occurrences of x must be subjectively the same, no such variation is possible. Certainly the definition 'sensation of a type which has such and such teaching links' does not allude to the subjective quality of the sensation, but only to the company it keeps. However, if it is a necessary condition (as I argue) of a sensation being of that type that it have a certain subjective quality, A. – then , 'the meaning must be private too' [*in a sense*, only].

The point of trying to separate meaning and reference is to show that C-subtle can be cleared of the charge that, [110] because it allows the possibility that the same name refers to subjectively different sensations in different people, so the meaning of sensation terms may be different for different people, and communication about sensations delusory. This charge could be dismissed if it could be shown that the meaning of sensation terms is no respecter of subjective variety. But in fact the charge sticks, once the requirement of intrapersonal subjective constancy is made. I say nothing here of more extreme attempts to separate meaning and reference, which issue in remarks like this one of Locke's (1968, p. 99): The rule "for a person to feel a pain a sensation must exist" is not a rule determining the meaning of "pain", but a rule determining when it is true to say of someone that he feels a pain.' I deal with this kind of approach in Hardy (1974), pp. 27–35. [Does 'meaning' = 'criteria for application?']

different in subjective quality. Indeed they can, and this situation is allowed by the extreme version of C-subtle discussed at the beginning of this chapter, according to which a sensation, to be an instance of x , needs only to occur in the right company: its **[111]** particular subjective quality (given that it is at least a sensation) is not in any way relevant. But the moderate version of C-subtle with which we are here concerned does make a requirement about subjective quality, i.e. that it remain the same for a given person. And this immediately rules out the possibility canvassed in the above objection, since once one has experienced x there is only one way for future occurrences of x to retain the same subjective quality. It is as if one was to require that once I had set eyes on Queen Elizabeth II, no other individual could ever count as queen of England for me. In that case it would become (however inappropriately) part of the meaning of ‘queen of England’ in my idiolect that Elizabeth Windsor be the person fulfilling the relevant role. And if everyone else was to mean the same by ‘queen of England’ as I did, they would have to be bound by this requirement too. Such a meaning of ‘queen of England’ is of course not the one the expression now has: but it serves to provide an analogy to C-subtle’s view of sensation terms, and so should clarify why this theory is committed to the interpersonal requirement that the subjective quality of x should be similar, and similar in a definite way, for everybody. C-subtle, in short, entails at least my version of ODM.

The third main argument against C-subtle points not only to the fact that the theory⁶⁷ sanctions the possibility of **[112]** an enormous gap between one person’s subjective world and that of another, a gap of proportions not contemplated in the context of localised discussions about particular sensations; but also, and more conclusively, to the impossibility of following C-subtle through to its logical conclusion.

⁶⁷ [Waiving the previous argument!]

The argument is in a way⁶⁸ an extension of the first one, going beyond both specific and general sensation terms to the term ‘sensation’ itself. But by being extended in this way it acquires a new force which it did not possess in its unextended form.

The argument is as follows. If C-subtle is to be consistent in its eschewing of interpersonal subjective requirements, it must give the same sort of analysis of the meaning of ‘sensation’ as it gives of a specific sensation term such as ‘*x*’. This analysis will have to be something like ‘thing (state? event? phenomenon? object?) of a type which (normally) has such and such teaching links’. The intrinsic quality of the ‘thing’ must, as in the case of a specific sensation or type of sensation, be constant for a given person, though no requirement of *interpersonal* comparability is made. My sensations need not, even in general, be like yours.

This development of C-subtle has two interesting consequences. First, it makes the connection of language with the inner world, which C-subtle (unlike behaviourism) does not wish to reject outright, perilously and unpremeditatedly slim, almost to the point of non-existence. If my version [113] of ODM is replaced by C-subtle, then although we have reluctantly to give up the idea that my sensation *x* is subjectively like your sensation *x* – or at any rate we have to accept that whether the two are subjectively similar has nothing to do with the meaning of ‘*x*’ – at least, we might think, both of us have a *sensation*. That seems a safe enough assumption. We know what sensations are like, and at least we both have one of those. But now even this solace is denied us. ‘Sensation’ is not to be a general term for a class of experiences whose essential, subjective nature we are acquainted with from our own case. Rather it is to mean, in effect, ‘whatever there is over and above any (set of) teaching link(s) which is/are definitive of a sensation’. Whether this ‘little extra something’ varies from person to person we are in no position to know: but such variation would not affect the meaning of the term ‘sensation’.

⁶⁸ [Only vaguely.]

This, once again, does not show that C-subtle is logically impossible: just that it does not achieve what at first sight it might appear to achieve. Its consequences are more far-reaching than they at first appear. All we are entitled to claim when we use the term ‘sensation’ is that something non-outer, of an unspecified and perhaps unspecifiable subjective nature, is in the offing, and is to be identified as a sensation purely in virtue of its (normal) observable accompaniments. What is left now to distinguish C-subtle from behaviourism? Not a nothing, [114] but not much of a something either.

But the second consequence of the definition of ‘sensation’ on C-subtle’s model makes even this picture of the meaning of ‘sensation’ impossible. In the case of a particular sensation term ‘ x ’, analysed as ‘sensation of a type which ...’, the meaning of ‘sensation’ is treated as given. This is all right. Given that we know what sensations are, we can learn that x is that sensation which is accompanied by such and such public events. We have a prior understanding of ‘sensation’ upon which to base our learning of the meaning of particular sensation terms. But is it the same with the general term ‘sensation’? Of what more general term do we have a prior understanding as a basis for learning the meaning of ‘sensation’? There seems to be no satisfactory answer. In the analysis given above I used the term ‘thing’: but this is a most (t,W,If\.,. $\epsilon \sim >$.. ($\$ \sim \sim$). unhelpful term (as are ‘state’, ‘event’, ‘phenomenon’ and ‘object’), little better than a plain variable. A sensation is not a thing in the most obvious narrower sense of the word, in which it designates physical objects. It can only be called an object or thing at all by a considerable stretching of these terms’ public meanings. Sensations are pretty much *sui generis*. So to call a sensation a ‘thing’ or ‘object’ is virtually to fail to characterise it at all (not that these terms are inappropriate when one already understands what sensations are). It seems impossible, in fact, without departing from C-subtle’s model, to learn the meaning of ‘sensation’.

One way to bring this out is to contrast the case of learning the meaning of ‘sensation’ with the situation in Wittgenstein’s

example of the beetles in their boxes. Wittgenstein uses this example to demonstrate the absurdity of conceiving of sensation language as I do, on the model of object and designation. In fact, though, because of the difference which I shall describe between this example and the case of ‘sensation’, it is not as strong a support for his position as he believes. (Of course, even if it *were* this strong, I would still not be persuaded to adopt his position.) And for the same reason it shows by contrast how peculiarly intractable the problem is of teaching the C-subtle meaning of ‘sensation’.

Wittgenstein writes:

Suppose everyone had a box with something in it: we call it a ‘beetle’. No one can look into anyone else’s box, and everyone says he knows what a beetle is only by looking at *his* beetle. – Here it would be quite possible for everyone to have something different in his box. One might even imagine such a thing constantly changing. – But suppose the word ‘beetle’ had a use in these people’s language? – If so it would not be used as the name of a thing. The thing in the box has no place in the language-game at all; not even as a *something*: for the box might even be empty. – No, one can ‘divide through’ by the thing in the box; it cancels out, whatever it is. (1953, §293)

Before I can use this example for the purpose in hand, I must first identify Wittgenstein’s mistake. He is wrong to say ‘The thing in the box has no place in the language-game at all; not even as a *something*: for the box might even be empty.’ For even if ‘beetle’ is **[116]** so defined that the intrinsic characteristics of individual beetles are irrelevant to their being beetles, we know none the less what general sort of item we are concerned with. We know what it is for something to be in a box, on the basis of our experience of cases exposed to public view. We know what sorts of things can be in boxes. So if the box is empty we shall know that it is: it is not true that this would make no difference. Even if the intrinsic qualities of beetles are linguistically neglected, ‘beetle’ can still be the name of a perfectly ordinary

sort of thing.⁶⁹ By the same token we could, if we wished, describe the contents of our boxes: there is no necessary bar to communication here. It is just that we have decided not to look in anyone else's box, and not to describe our own beetle. If this convention were dropped, beetles could be described and compared like anything else. Talk merely of 'beetles' is not the best that language can do (though it is possible to imagine circumstances in which it may be the most convenient way of talking).

The same sorts of things cannot be said of 'sensation', as conceived by C-subtle. It might be claimed that just as we could learn the meaning of 'beetle', in the imaginary case, as 'thing in the box', so we could learn the meaning of 'sensation' as 'thing with such and such teaching links': that 'beetle' in Wittgenstein's example has a meaning analogous to that of 'sensation' as conceived by C-subtle. But in fact we could not supplant link talk with talk of intrinsic sensation quality, as we can supplant [117] beetle talk with talk of the intrinsic properties of beetles. There is no description we can give of what sensations in general are like intrinsically. We cannot appeal, moreover, as we can in the case of 'beetle', to an already understood notion of material object, or to an already understood notion of one thing being in another, in order to give some clue as to what sort of item a sensation is supposed to be. All we can say is that there is sometimes 'something' present in given company, something which is not publicly observable, and not in the least like any other sort of thing. And this is quite uninformative, by itself.

The fact is that any genuine understanding we have of the meaning of 'sensation' must come from our own personal experience of sensations, and cannot be provided by purely external tuition. When teaching the meaning of 'sensation', we have to rely on the assumption that our pupils have sensations like ours, and can grasp that it is these that we are referring to when we use the term 'sensation'. We cannot characterise

⁶⁹ [Or 'perfectly ordinary things'.]

sensations in terms of other previously mastered notions drawn from that part of language which deals with public phenomena. In other words, unless we accept the argument from analogy, we cannot explain how we come to understand what sensations are.⁷⁰

[118]

If, then, there is no more general public concept to which appeal can be made in learning the meaning of ‘sensation’, it seems that C-subtle’s programme cannot even get started. The whole enterprise is radically misconceived. Sensation language cannot be taught, as a whole, in the roundabout way imagined by C-subtle. In which case, once again, C-subtle’s attractiveness as an analysis of the meaning of individual sensation terms is greatly diminished.

So we are left with the choice between behaviourism,⁷¹ my version of ODM, and any other alternatives there may be. How are we to choose between rivals that are, logically, equally possible, in that they can all [119] encompass the observed facts of linguistic behaviour? In the last analysis, perhaps by making a

⁷⁰ It may be objected that I am confusing a thesis about the meaning of sensation terms, which is what C-subtle is, [118] with a thesis about how meanings are learnt, which C-subtle is not. Could it not be the case both that C-subtle was true, and that the meanings of sensation terms were learnt through personal experience? For example, even if ‘sensation’ means ‘thing with such and such teaching links’, could this not come to be understood by a pupil on the basis of his having a sensation, noticing the concurrent presence of the teaching links, and saying ‘Ah, *now* I see what you mean?’ The answer to this is that this would be possible, though it is against the spirit (if not the letter) of C-subtle. The motive for adopting C-subtle is to avoid the charge that sensation terms might mean different things to different people, and this is done by analysing them in such a way that, supposedly, subjective experience does not contribute to the understanding of their meaning. But the learning procedure just imagined does crucially implicate subjective experience, without which the C-subtle analysis of ‘sensation’ could not come to be understood. And if the only way to teach C-subtle meanings is to rely on the very assumptions whose insecurity provided the grounds for adopting C-subtle, there is no longer any reason (of that kind at any rate) to persist with C-subtle.

⁷¹ [Not really.]

C-SUBTLE

decision. But there are reasons for deciding one way rather than another, some of which will emerge in the next chapter.

3 WHAT DO SENSATION TERMS MEAN?

Some philosophers think that something's having intuitive content is very inconclusive evidence in favour of it. I think it is very heavy evidence in favour of anything, myself. I really don't know in a way what more conclusive evidence one can have about anything, ultimately speaking.

Kripke (1972), pp. 265–6

[T]he meanings of words are affected, and often very deeply affected, by our explicit or implicit causal beliefs, and the analysis of what is meant by an expression may very well reveal all kinds of physical or social or psychological beliefs or assumptions prevalent in a given society, a change in which could affect the meanings of words.

Berlin (1950), p. 311

I return now to the question with which I began the first chapter of this thesis: What do sensation terms mean? Numerous hints, and more than hints, have been given, apropos of other issues, as to the sort of answer I would give to this question. In this final chapter I aim to develop these hints; to confront the question directly.

The question can be taken both in a general sense and in a particular sense. If it is taken in the general sense, what is being asked for is an all-inclusive account of the meaning of all sensation terms, considered together as a species of term – an analysis that will [121] apply to any sensation term, whatever its own special, unshared characteristics may be.⁷² It cannot, of course, be a foregone conclusion that such a general account is

⁷² [‘What is it to be a sensation?’]

WHAT DO SENSATION TERMS MEAN?

possible: it may be that the meanings of individual sensation terms are too various to be comprehended by any single all-embracing formula. Though there may be a strong presumption that some degree of generalisation will be possible, if only because of the existence of the general description 'sensation term', it should be shown what this generalisation consists in, not assumed that it is obvious.

If the question is taken in the particular sense, it would be answered by a list of the meanings of all sensation terms, or at least of a large and representative sample of these terms. Each item in the list would be independent, and no generalising hypothesis, to link them all together, would be required – though, once again, it would be only natural, after reading such a list, to formulate such a hypothesis, or at least to ask why all the items in it should be thus grouped together under the one heading as 'sensation terms'.

In what follows I shall make some remarks in answer to both versions of the question, though my main interest is in the first, general, sense, so that what I say about the meanings of particular sensation terms will be directed by and large towards discovering or justifying a general answer.

If the arguments put forward in chapters 1 and 2 [122] are accepted, all the options, apart from holding C-subtle to be true across the board,⁷³ remain open to us in our attempt to discover the meaning of sensation terms. The whole range of answers that seem, on first looking into the matter, to be possible is still available. What are these answers? At one extreme of the range is the view, already mentioned, which Pears calls 'C-crude': the view that sensation terms have what might be called an entirely 'intrinsic' meaning – that a sensation is the sensation it is solely in virtue of its subjective quality.⁷⁴ At the other extreme (if we discount behaviourism, which has nothing properly to do with

⁷³ As stated in chapter 2, the option does still remain open of giving a C-subtle type of analysis of the meaning of any individual sensation term.

⁷⁴ [In which case links are mere reference-fixers.]

WHAT DO SENSATION TERMS MEAN?

sensations) is the version of C-subtle whose unsatisfactoriness as a general theory I attempted to demonstrate in the last chapter – though, once again, it is still available for use in particular cases. This theory favours what we might call an ‘extrinsic’⁷⁵ meaning for sensation terms, whereby a sensation would be identified as being of a certain type by the publicly observable context in which it occurs, roughly speaking: different sorts of public factor, as we shall see, would [123] be relevant for different sensations, but this variety does not affect the general point that meanings of this kind would be ‘extrinsic’ rather than ‘intrinsic’.

Between these two extremes there is any number of intermediate positions. But they would all require in common that both intrinsic and extrinsic conditions should be satisfied by a sensation if it was to count as a sensation of a certain kind. More or less emphasis might be placed, in an individual case, on intrinsic or on extrinsic requirements, but so long as the two extremes mentioned above were avoided, both intrinsic and extrinsic properties would remain jointly necessary (and sufficient) for determining the identity of a sensation. This view is the version of ODM I adopted provisionally at the outset, the view I have had in mind when defending the possibility of making definite intrinsic requirements of a sensation before allowing it to qualify for a particular name. I shall be coming back to it below.

Where in this range do the meanings of our actual sensation terms lie? At one of the extremes? Somewhere in between (and if so, where exactly?)? Is the answer the same for all sensation terms, or even for all uses of a given sensation term? And what method (this question should really be answered first) should we adopt to discover the meaning of a sensation term?

To answer such questions exhaustively is beyond the scope of this thesis. My main purpose – by now, I hope, achieved – has been the negative one of warding off certain [124] a priori restrictions on the discussion out of which answers may come.

⁷⁵ *OED* sense 4a: ‘pertaining to an object in its external relations’.

WHAT DO SENSATION TERMS MEAN?

But of course this negative achievement is of value only in so far as it better enables the positive enquiry to proceed in an unencumbered way, and ultimately only in so far as the upshot of that positive enquiry is as a result nearer the truth than it would otherwise have been. Naturally I believe that these consequences will follow, because I believe that subjective experience does have an important part to play in an account of the meaning of sensation terms. And so, although I shall not go into the question of the meaning of sensation terms with the detail it deserves, I am bound to make some prefatory remarks, at least, on the subject, and these are what this chapter contains. The remarks will be often preliminary, programmatic and eclectic in character: my object is to provide a framework within which the enquiry can be elaborated, rather than to attempt a definitive account of such a complex field.

A framework of enquiry

I begin by making one or two very general distinctions, and raising a number of questions which arise out of them, in order to provide some further points of reference around which to organise subsequent discussion.

So far, in considering possible meanings of sensation terms, I have made only the broad distinction between intrinsic and extrinsic factors. And indeed, since it [125] is my main concern to rescue the reputation of intrinsic factors, this relatively crude distinction has served well enough up to this point. But when we turn to the detailed spelling out of the meanings of particular sensation terms, we need a finer mesh through which to sieve our data. Both categories, 'intrinsic' and 'extrinsic', must be subdivided. Such subdivision could be proliferated endlessly, to the point where each sensation term had a category all to itself – a *reductio ad absurdum* of the process of refining a system of classification. For present purposes it will be more useful to subdivide only once or twice.

WHAT DO SENSATION TERMS MEAN?

First, then, let us subdivide the extrinsic factors into those which characteristically precede the occurrence of a sensation, those which occur at the same time, and those which follow it. In order to have a short way of referring to each of these types of factor, we could call them respectively ‘pre-extrinsic’, ‘co-extrinsic’ and ‘post-extrinsic’; and these labels could also be used to describe sensation terms which turn out to have meanings in which one or other of these types of factor predominates. For example, if the meaning of ‘itch’ has mainly to do with a desire to scratch, ‘itch’ is a co-extrinsic term; and if the meaning of ‘pain’ is largely a matter of the cause of the sensation,⁷⁶ ‘pain’ is a pre-extrinsic concept. These categories might be further subdivided into [126] ‘causal’ and ‘non-causal’, according to whether the extrinsic factors that feature in an analysis of the meaning of a term stand in a direct causal relationship to the sensation in question. This division would yield categories such as ‘causal pre-extrinsic’, ‘non-causal post-extrinsic’ and so forth. Perhaps ‘stiffness’ is partly a causal pre-extrinsic notion, and ‘sleepiness’ a non-causal post-extrinsic one: characteristically, stiffness is caused by exercise, and sleepiness is followed by, but does not cause, sleep.

The intrinsic category subdivides in a different fashion, such that we cannot separate one intrinsic subcategory from another in the way that we can in the case of the extrinsic category. It is possible, for example, for an extrinsic factor to precede a sensation without following it (as in the case of electric shock); but a sensation must always have both a quality and an intensity – one without the other is inconceivable (though it may be that one can vary while the other is held constant). So the subcategories of the intrinsic category will correspond to different aspects of subjective experience, not to different types of experience capable of independent existence. Such potentially independent types of experience do of course exist. For a start there are bodily sensations, and the five different types of experience brought to

⁷⁶ [But it isn’t?]

WHAT DO SENSATION TERMS MEAN?

us by the five senses, which despite their important interactions have a degree of independence, as is shown by the cases of people who lack one or more of [127] them; and no doubt these five types could be further subdivided along similar lines (e.g. touch into sensations of texture, temperature and shape). But the deliverances of the five senses can also be distinguished extrinsically, by the kind of information provided, by type of stimulus, by mediating bodily organ;⁷⁷ and besides, we have the intrinsic subcategory ‘quality’⁷⁸ to cover any subjective mirror of the extrinsic division of sensations by sense. ‘Quality’, then, is one intrinsic subcategory, ‘intensity’ another; and to these should be added ‘temporal pattern’. And, just as in the case of extrinsic categories, we can use these intrinsic categories as a basis for naming the appropriate type of sensation term: if ‘sensation of red’ is called ‘quality-intrinsic’, this means that it is quality of subjective experience that critically determines that a sensation is a sensation of red; if ‘agony’ is ‘intensity-intrinsic’, the distinguishing criterion of agony will be its subjective intensity. (All these examples, of course, though not picked at random, are offered only *exempli gratia*, to illustrate the categories: I do not necessarily believe that the terms I mention are exclusively of the kind to which I hypothetically assign them.)

Extrinsic factors too, of course, can also vary in temporal pattern, and in intensity (in a different, objectively defined sense of ‘intensity’): and it may be that there are some sensations which are picked out by variation along these parameters – perhaps, for example, [128] part of what differentiates a shudder from a shiver might be the intensity of bodily vibration.

Given an apparatus of categories such as this, we can examine the meanings of sensation terms and see how the categories are deployed. Does every sensation term involve the same categories, or one of a limited set of category combinations? Do any

⁷⁷ [Not for touch.]

⁷⁸ [Is quality (partly) determined by extrinsic factors? If so, how? This is mysterious.]

WHAT DO SENSATION TERMS MEAN?

sensation terms involve only a single category, or is there always more than one in question? Is the typology of categories enlightening, either in part or as a whole, or does it yield only an uninformative allocation of sensation terms to unexplanatory classes? Are some categories empty? Can we make any suggestive generalisations, such as that every sensation term will have at least one intrinsic and one extrinsic element in its meaning? If the meanings of sensation terms vary in structure from case to case, is there any discernible pattern in this variation? Can we say, for example, that generic sensation terms are more extrinsically oriented? – is it more important that the perceived quality of red should have a certain definite subjective character than that ‘colour sensation’, or, to be even less specific, ‘visual impression’, should always and for everybody cover the same range of subjective experience? What differentiates closely related types of sensation such as giddiness and nausea, or aches, pains and cramp?

These are examples of the questions that suggest themselves. When they were answered, they would in turn [129] suggest further questions about why the meanings of sensation terms are as they are. Why is it (if indeed this is true) that some terms hook on predominantly to the pre-extrinsic end, and some to the post-extrinsic end, of the characteristic input–sensation–output series of events that provides the normal environment for the occurrence of a particular type of sensation? Why do other terms (if indeed there are any such) apply rather on the strength of the intrinsic nature of the sensation? In general, what determines which of our categories a sensation term falls into? Pedagogic constraints? Considerations of utility? Some third thing?

It should be clear by now how vast is the field of enquiry opened up by the question ‘What do sensation terms mean?’

WHAT DO SENSATION TERMS MEAN?

Methods of enquiry

Before any of these questions, descriptive or explanatory,⁷⁹ can be answered, something must be said in reply to the logically prior question, already mentioned, ‘What method should we adopt to discover the meaning of a sensation term?’ Until we have an acceptably secure procedure to hand, the meanings we come up with may be little better than guesswork, and any answers or generalisations based on them would be correspondingly unreliable. Maybe there is no foolproof method available, in which case we must use the best there is. What are the candidates? **[130]**

Concepts evolve partly to fulfil purposes – to pick out what we need to pick out, to make useful distinctions, to enable us to refer to and describe those elements and structures in our environment which is in our interests, for one reason or another, to classify together and converse about. This, if vaguely put, is at least, I hope, uncontentious. So one basic strategy to adopt in discovering the meaning of a term must be to ask what work it does for us. This is likely to reveal part, at least, of its meaning – not necessarily the whole of it, since, as we shall see, the functional⁸⁰ role of a concept is not always the only factor that determines its meaning. **[131]**

⁷⁹ [i.e. questions of description or explanation.]

⁸⁰ Let me ward off a possible ambiguity of the term ‘functional’. It is possible to use the expression ‘functional meaning’ to pick out concepts definable in terms of the function performed by what falls under them. If ‘chair’ means ‘thing for sitting on’, it has, in this sense, a functional meaning; so does ‘pain’ if it means ‘sensation that indicates injury’; whereas a natural kind term: such as ‘water’ has, if any meaning at all, a non-functional one. But I do not use ‘functional’ in this sense. The function of a concept is to be understood as the purpose it serves in our lives: so ‘chair’ is a functional concept because it enables us to pick out just those things which satisfy certain broad structural conditions and can be sat on; ‘pain’ has the function of enabling us to refer economically to sensations that (perhaps among other things) signal injury; and so forth. In this sense of ‘functional’ a complex

WHAT DO SENSATION TERMS MEAN?

It is by no means always self-evident exactly what work a concept is doing (if it were, there would be little occupation for philosophers); and some sensation terms, especially perhaps ‘pain’, seem to be particularly elusive in this respect. So we could do with some help in the search for their meanings. This help, I suggest, is forthcoming from at least two important sources.

First, an examination of the teaching links⁸¹ of a sensation term – the public phenomena whose co-occurrence or near co-occurrence with types of subjective experience makes it possible to teach the meanings of sensation terms to language learners – may also reveal the function the term performs, for it will be natural for function and teaching links to overlap, as we shall see when we come to the consideration of specific examples. This approach is helpful too in cases where a meaning is not entirely functionally determined – where there are elements present in the meaning which would not be justifiable in a language constructed in accordance with some rigidly utilitarian ideal. How this is so we shall see, again, when we come on to consider specific examples.

As a second strategy we may use the familiar device of asking what we would say if the nature of things [132] were to change in certain ways. To ask what we would say if things were different is a good way of discovering what we mean when we speak of things as they are. It can reveal assumptions which are so thoroughly taken for granted as to be easily forgotten, even though they may be all-pervasive. It can show what are the crucial factors governing the use of a concept, and what are peripheral considerations whose absence would not be much noticed. And this, once again, can reveal both the functional and the non-functional elements in a term’s meaning.

Let us look at each of these strategies in rather more detail.

concept can have both functional and non-functional elements – elements that can and elements that cannot be justified in rigorously utilitarian terms. (Language games have both purposive and non-purposive rules.)

⁸¹ Cf. p. 59, note.

WHAT DO SENSATION TERMS MEAN?

Teaching links

It is necessary to emphasise the variety of teaching links.

Pears (1971), p. 149

So far I have concentrated on only one or two sensations (apart from a handful mentioned in the preliminary remarks above, by way of illustration), and so on a narrow range of teaching links. It is time now to expand this range somewhat, in order to have a more representative sample of data on which to base any generalisations about the connection between teaching links and meaning.

First, then, we must introduce some more sensations. It is notable that it is difficult to think of a list of sensation terms more than two or three items long, unless one includes those sensations that can only be described [133] by borrowing the name of a publicly observable event, object or other phenomenon which is in some way closely associated with them, by being a cause, or an effect, or an expression, or by having some less direct connection which can nevertheless be used to underpin the use of the name of the public item also as (part of) the name of the sensation. This is how the vast majority of sensation talk is conducted. In a few cases, notably that of 'pain', there is a special word in the language for a type of sensation. These cases may have something to do with the frequency of our need to mention certain sensations; or with their peculiar importance to us; or with the fact that some sensations have such a multiplicity of causes and/or effects that there is no obvious single public item or class of items which could plausibly lend its name to them. At any rate, many of the items in the following catalogue of sensations should strictly speaking be preceded by the expression 'sensations of': by themselves they are not really names of sensations at all. But since we do commonly use the name of a sensation's cause (for example) as a shorthand name

WHAT DO SENSATION TERMS MEAN?

for the sensation itself, in this catalogue I follow and extend this licence, if only for the sake of brevity.

There are, then, pains, itches, tickles, aches,⁸² headaches, stomach aches, agony, discomfort, giddiness, [134] cramp, tinglings, throbbings, nausea, stiffness, sleepiness, fatigue, migraine, glowings, tension, lethargy, inertia, palpitations, shivering, hypersensitivity, shock (electric and otherwise), myriad tastes and sounds and colours, cold, heat, pins and needles, sensations of shape and texture, of tone and timbre, of wanting to yawn or laugh, of vomiting, of illness, of being out of condition, of having nudged the funny bone from a sensitive angle, of scraping a piece of chalk or a fingernail along the surface of a blackboard, of blushing and of sweating, of wind on the flesh, of orgasm, sexual arousal, indigestion, sneezing, a full bladder, hunger, thirst, emotional disturbance (fear, anger, hate, love etc.), the ‘jumping’ sensation while falling asleep, and so on, almost endlessly.

What are the teaching links of these sensations? In many cases the answer is self-evident: the name of the sensation is a giveaway. The public occurrences by which one is enabled to teach and understand the meaning of expressions like ‘electric shock’, ‘sensation of orgasm’, ‘sensation of thirst’ are too obvious to need spelling out. And equally, to the extent that the function of a term is exclusively to pick out just those sensations characterised by having the public accompaniments used as its teaching links, it will be obvious what the meaning of the term is – indeed the term itself will be almost its own definition: ‘the “jumping” sensation while falling asleep’ might be an example. Whether there are any such terms, definable exclusively in terms of their links, [135] In C-subtle fashion, is something I shall be

⁸² It does not seem to me that either headaches or stomach aches are species of the genus *aches*, if by ‘aches’ is meant muscular and bone aches. Certainly there must be a reason why these sensations are all called aches, but to some extent, too, they do seem to be mere homonyms.

WHAT DO SENSATION TERMS MEAN?

considering in due course. First, more needs to be said about the identifying of teaching links in less straightforward cases.

In the case of many descriptions of sensations it is not so immediately obvious what the teaching links are, or how they operate, for there is no superficial verbal evidence to hand – or none that is not misleading. I shall discuss one or two cases, without pretending to be comprehensive.

Let us begin with the natural expression of pain. When people are in pain they tend, if the pain is severe, or if their tempers are frayed, or if they are young or unselfcontrolled, to ‘wince, scream, flinch, sob, grit their teeth, clench their fists, exhibit beads of sweat’⁸³ (Putnam 1965, p. 9), and so forth. Assuming that these responses are innate (and surely at least some of them are), we can say that they provide natural links for teaching the use of ‘pain’.⁸⁴ **[136]**

But, as Pears points out, there are refinements in pain talk which cannot be learnt on this simple model. Here the cause of pain, or its pattern, as opposed to its expression or effects, can be useful as links, particularly in learning to describe the quality of a pain. Pains can be throbbing, shooting, burning, stabbing, prickling, searing, scorching, chilling and so forth (‘acute’ in some uses probably comes in here too). Descriptions of this kind can be subdivided into two groups. The first group describe the pain by analogy, by implying that it has a quality like the pain one gets when one is really being burnt, pricked, chilled etc. The second group draws attention to the pain’s pattern: a throbbing pain has a certain temporal intensity pattern; a shooting pain travels a fair

⁸³ Beads of sweat are not an ‘expression’ of pain: rather they are an effect (and so a sign). How is one to distinguish expression and effect? The distinction is linked to that between actions and passive reactions. Wincing is something one does, albeit involuntarily; sweating is something that happens to one. But the distinction is not important, in this context, for the language pupil.

⁸⁴ Perhaps we should add ‘expressive co-/post-extrinsic’ to our list of categories, in case there turn out to be terms with meanings of the kind ‘sensation with such and such an expression’.

WHAT DO SENSATION TERMS MEAN?

distance in the body at high speed, and is reasonably intense besides; a stabbing pain is sharp and sudden and precisely located. It is not always easy to decide which group a description belongs to: is 'sharp' used by analogy with pain received from sharp objects, or is it a reference to intensity, and narrow and definite location? What sort of analogy (if any) is 'dull'? Also, it is sometimes easy to make a mistake: 'stabbing' might appear to suggest similarity to the pain of being actually stabbed – but how many of us have ever experienced the penetration of our flesh at high speed by a long, sharp, metallic object?

Strangely, when Pears talks (1971, p. 148) of the sort of distinction which can be taught with reference [137] to the type of stimulus which causes the pain, he does not give examples of the kind belonging to my first group, analogical descriptions. Rather he mentions the distinction between the surface and interior of the body, and between interior bodily locations. But how can an appeal to stimuli help us here? Surely this kind of distinction can be made in just the way Pears in fact here makes it, viz. by reference to the precise part of the body where the pain is felt (the language pupil may be supposed to have mastered the description of bodily location already).

Other distinctions rely on the fact that certain types of subjective experience occur in company with certain public events (either before, simultaneously with or after them), and yet are neither their causes, nor their effects, nor expressed by them. An example of this kind, mentioned by Pears, is the sensation that normally immediately precedes a yawn: it is not the cause (or effect) of the yawn (though the sensation and the yawn presumably share a cause), nor is the yawn its natural expression; but it can safely be identified as a yawn's predecessor. A similar case is the distinction between pain and nausea: nausea is, or involves, an unpleasant sensation which may be accompanied by all the normal symptoms of pain mentioned earlier, and yet it arguably isn't a kind of pain. At any rate, it can be distinguished by being the (non-causal, though again presumably co-caused)

WHAT DO SENSATION TERMS MEAN?

harbinger of vomiting, an impeccably public event.⁸⁵ Also in this class is that [138] sudden muscular jump one's body gives as one is dropping off to sleep: unless this experience occurred regularly at this or at some other publicly identifiable point in our daily lives, it would be much harder to communicate about it without our being in doubt as to whether we were all talking about the same experience. It would not be sufficient to characterise it as the experience of a muscular spasm: there are spasms and spasms, and this one (it seems to me) feels special.

For the sake of completeness, though it will not impinge on our subsequent discussion of the relation between teaching links and meaning, I should mention finally what Pears calls 'the ways in which the pupil fills in the gaps, and makes moves which go beyond his literal instructions'. These lateral moves may be vital, as in the case of descriptions of sensations like 'pins and needles', which we can understand without direct experience of the origin of the analogy: there is no other kind of way to describe many subjectively distinguishable experiences. Other such moves may simply provide time-saving short cuts: Pears mentions the transference of the concept of intensity⁸⁶ from one field of sensations to another. The intensity of a [139] pain is not like the intensity of a sensation of lethargy, still less of a sensation of red or of the taste of beetroot, beyond the basic fact that there is 'more' of any intense sensation than there is of its mild counterpart. But just as we can transfer the concept 'more' to subject matters widely different from those in connection with which we learnt its use, so we can transfer the concept of intensity easily enough from one sensation to another, although it would be possible to learn it separately for each.

⁸⁵ [Also by inner quality!]

⁸⁶ What is it for a sensation to be intense? Is intensity to be measured *intrinsically* – and if so, should the scale be interpersonal, or relative to the relevant individual? – or *extrinsically*, in terms of the energy of the relevant public symptoms, say? All three possibilities would seem to make their contribution to the way we actually talk about the intensity of sensations.

WHAT DO SENSATION TERMS MEAN?

Such in very broad outline is the variety of teaching links used to teach the use of sensation language. There will be many more refinements; each sensation term will be liable to have its own special features. Wittgenstein's too exclusive concentration on the natural expression of pain led him, as Pears puts it, to treat 'what are really descriptions of sensations as if they were verbal expressions of sensations' (cf. p. 10, note). Perhaps we have now accumulated a sufficiently various collection of links to provide a basis for enquiring how knowing what the teaching links for a sensation term are can help us to discover its meaning – and, maybe, how it cannot help us.

The fact that certain teaching links are used to teach the meaning of a sensation term does not, of course, by itself entail that it is part of the meaning of that sensation term that such links are present whenever the [140] sensation in question is experienced. All that is established is a pedagogic role for the links, not a semantic role. Some teaching links may feature in meaning analyses, others not. But there is at least a strong presumption that teaching links are semantically implicated, for one of two reasons. The first reason, already alluded to, is that the point of the term may be connected directly with the links, in that we pick out the type of sensation in question partly or wholly because of the public accompaniments which also serve as teaching links: unsurprisingly, observable phenomena which motivate us to isolate linguistically a particular [subjective] type of sensation often also provide a convenient route to understanding the name by means of which we isolate it. For example, the obvious link for 'itch' is scratching, and since presumably one of the main points of having the term 'itch' in our language is to be able to pick out, and so tell people about, sensations that make us (want to) scratch, here is a case where link reveals meaning in the way described. We learn the meaning of 'itch' via scratching; and 'itch' means 'sensation that makes you scratch'. This may be a very simple case, and perhaps there is no need to go via the teaching link stage to discover that scratching has something to do with itching, and so with the meaning of 'itch'. But it shows in

WHAT DO SENSATION TERMS MEAN?

a clear way the sort of connection which, in a more obscure instance, might be revealing. Certainly the insights [141] about 'pain' which Wittgenstein suggested by drawing attention to the way in which the concept is taught are not trivial or self-evident,⁸⁷ at least not without the benefit of hindsight.

The other reason for the presumption that teaching links are semantically implicated is this: regularly present links will tend to be involved in meaning even if no clear function is served by their doing so, simply because they dog the heels of the sensation whose name they help to teach.

It may be that they actually encumber the concept, by saddling it with extra necessary conditions of applicability which obstruct the function which it might be held primarily to fulfil. It may be worth saying a little more about this, since it is a phenomenon that occurs elsewhere in language, not merely in the case of sensation terms. Whenever the [142] functionally inessential properties of a kind of item are particularly regularly present, perhaps as regularly as the functionally essential properties, then they tend to be so entrenched in the collective mind that their presence is also a semantically essential requirement.

This phenomenon can be illustrated clearly by using examples of terms for man-made objects. Here the inessential properties can actually be watched over time, growing parasitically on their initially unsullied hosts.

First, processes of standardisation change a word's meaning in this way. I quote from Kovesi (1967), p. 14:

⁸⁷ To the extent that the meaning of 'pain' is pre-extrinsic rather than post-extrinsic (see further p. 119), there will be this difference between 'itch' and 'pain', that the semantically relevant teaching link for 'itch' is alleviative behaviour, while that for 'pain' is injurious stimulation. But the analogy relevant for the present context remains: that methods of teaching and elements of meaning overlap. One cannot rely on this overlap, as I say below: indeed Wittgenstein's belief that pain's expression is important for the meaning of 'pain' may be mistaken, for all that I use him in the text as an example of the efficacy of approaching meaning via links. None the less, a fallible guide is not no guide.

WHAT DO SENSATION TERMS MEAN?

Before people thought of manufacturing levers one didn't go to a special shop to get them but one asked for a lever from anyone likely to have the sort of object that would do the job. What was or was not a lever then depended solely on what did or did not fulfil the function of lifting objects in a particular manner. But once levers came on the market the situation was different. If one asks in a hardware shop for a lever now, the assistant cannot go out into his backyard and look for any piece of metal that would do the job. Manufacturing, buying and selling introduced new criteria for what will or will not be accepted as levers.

What used to be incidental features of some levers (that they had a certain rather precise shape, and were made of one of a rather limited range of metals) have now become defining properties: without these characteristics an object might still be *used as* a lever, but it would hardly, except in the context of technical mechanics, be a lever. Whether we find this particular example convincing or not, the point it has to make is clear enough.

Consider next the new process of three-dimensional [143] photography called 'holography', in which a three-dimensional image of an object is formed in mid air by the intersection of special laser beams. It is not strictly speaking (as the etymology reflects) part of the meaning of 'photograph' that the image be printed on paper or projected on a screen: but if we hesitate to call a holograph a photograph, this may reflect the fact that our concept of a photograph has become infected by our taking for granted what is really only one way of presenting a photographic image.

Imagine finally that King's College Cambridge was bombed to rubble, and a new building erected in another part of Cambridge to take its place, while the ruins of the old buildings were preserved as a memorial, after the manner of the shell of the original Coventry Cathedral. Might we not have some hesitation about whether the identity of the college was intact? This would not be simply because 'King's College' refers ambiguously to a set of buildings and a society of scholars, but because an institution

WHAT DO SENSATION TERMS MEAN?

becomes associated with its particular (inessential) architectural home in such a way that the survival of the home becomes to some degree definitive of the survival of the institution.

To return to sensation terms: here we are not (usually) concerned with a process through time of the kind just illustrated. There were not once upon a time pure sensation concepts, exclusively utilitarian, and then a process of corruption, of infiltration by alien criteria, [144] leaving the concepts cluttered with unserviceable paraphernalia. (This is not to say that sensation concepts sprung fully formed from the mind of prehistoric man: but their evolution was surely not of *this* kind.) Rather such concepts have always contained some elements which are more easily justifiable in functional terms than others. In one or two special cases, such as when new terms are deliberately invented for special purposes, a temporal debasement may occur. But normally the infection is genetic rather than environmental.

Nevertheless, the analogy with the examples discussed above is instructive. Even if certain public phenomena are made use of, primarily, merely as useful aids to teaching the meaning of a sensation term, their dependable presence has given them a secure semantic role. This is especially likely to be true of sensation terms whose function is to have a predominantly intrinsic meaning. If we want to communicate with other people about subjective experience, then because this experience is private we have to fall back on its publicly observable causes, effects, symptoms, manifestations, expressions or whatever in order to get our message across. These links do not need to feature in the meaning of sensation words used for this purpose, but since, nature being regular, they are reliable signs of the presence of the types of subjective experience with which they are correlated, it is not surprising if there is a strong psychological [145] connection between the links and the types of experience, such that we would be nonplussed to discover that all the characteristic links were present, but no experience to match; or

WHAT DO SENSATION TERMS MEAN?

that an experience of a certain kind was present, but no characteristic links.

This psychological connection is often logical too, and the presence of the links is a necessary condition for the applicability of the description in question. Sometimes it will be hard to decide whether the connection is sufficiently entrenched to affect meaning in this way. Sometimes it will be clear that it is. Many would claim, plausibly, that the role of the concept 'pain' in our lives is such that at least its pre-extrinsic links are undoubtedly important to its meaning. But often things are less clear; the precise boundaries of a concept, if indeed it has any, tend to be hard to pinpoint; and moreover, meanings are always to some extent in a state of flux, particularly under the influence of scientific discoveries. So which links are involved in meaning may be changeable as well as unclear at any given moment. The division between card-carrying members of the semantic union and mere fellow-travellers will be blurred, and there will be traffic across it. But for all that, the phenomenon of semantic participation by functional inessentials⁸⁸ does occur, and does provide one explanation of the connection between teaching links and meaning.

For one or other of these reasons, then, teaching links are liable to be a good guide to meaning. Even [146] so, it must not be automatically assumed that they are semantically significant, since it always remains possible that a useful pedagogic aid (perhaps because it is not universally applicable) is quite unconnected with the meaning of the term it helps to teach. To take a simple example, a child may learn the meaning of 'red' by being told that it is the colour of pillar boxes, ripe tomatoes, blood and so forth: but of course it is no part of the meaning of 'red' that any of these things are red.

There is, however, a far more serious shortcoming of the links method of investigating the meaning of sensation terms. It tells us nothing about the importance (or unimportance) of intrinsic

⁸⁸ [e.g. pain's expression?]

WHAT DO SENSATION TERMS MEAN?

factors. Of its nature it can be informative only about extrinsic elements in the meaning of sensation terms, since teaching links are necessarily extrinsic to subjective experience. We can communicate about intrinsic aspects of sensations only by means of public criteria, and so, if we read off the meaning of sensation terms solely from the links used in teaching them, from the criteria used in applying them, we will be driven to the conclusion, together with many before us, that intrinsic factors have no part to play.

This, though, is exactly the conclusion I have been trying to avoid all along: at least, I have been trying to avoid reaching it for this sort of methodological reason – that is, because there are no public criteria of the quality of subjective experience. Intrinsic factors, [147] I have argued, can feature in the meaning of sensation terms.⁸⁹ As with extrinsic terms, this can happen in two

⁸⁹ In case it is maintained that, because subjective quality cannot be directly described in public terms, it cannot be directly involved in meaning in the way I claim, but can at best be only the referent of a sensation term, it may be worth explaining how, even so, subjective quality can be crucial, in rather a different way. Wittgenstein believes that what has no public manifestation has no [linguistic] significance. He says (1953, §304): “But you will surely admit that there is a difference between pain-behaviour accompanied by pain and pain-behaviour without any pain?” – Admit it? What greater difference could there be? – “And yet you again and again reach the conclusion that the sensation itself is a *nothing*.” Not at all. It is not a *something*, but not a *nothing* either! The conclusion was only that a nothing would serve just as well as a something about which nothing could be said.’ But that a nothing will serve just as well as a something is true, on this view, *only from the point of view of the publicly observable requirements of language*. It is a mistake to generalise from what will serve for language to what will serve, period. This is like arguing that since from the point of view of the camera it makes no difference whether the racing car is made of steel or of papier mâché – because the photograph will contain no information about what lies beneath the paint – it simply doesn’t matter at all what the car is made of. But this is absurd: if the car was not made of metal it would not be tearing around the track for the camera to photograph. Similarly in the case of sensations: as far as the public criteria for learning and regulating the use of sensation language go, it may be irrelevant what quality of experience accompanies behaviour. But this does not mean it is simply irrelevant. Even if language is no visible respecter of subjective

experience, the nature of such experience still matters to us. To put this another way: even if we have to talk about sensations in such a way that it makes no logical difference what their subjective qualities are (of course, I do not believe that this is so), this need not mean that it is not often because of their intrinsic nature that we are interested in talking of them at all. For instance, we accord sympathy to someone in pain not because he is wincing and groaning, but because we believe he has the kind [148] of sensation to which these symptoms point: we know what sensation goes with similar symptoms in our own case, and we know how unpleasant it is.

Compare the situation, once more, in Wittgenstein's example of the beetles in the boxes. Here it can matter in a parallel way whether the beetles are the same. Certainly, it is logically possible for everyone to have something different in his box, given the way the example is set up. It is also true that whether or not this was the case would not affect the logical possibility of the word 'beetle' having a use in the language. But it may remain true that, logical possibilities notwithstanding, the reason these people talk about beetles (the reason we talk about pains) is that they want to talk about them qua having the intrinsic properties they have. (If this is their aim, it is a trifle perverse of them to go about achieving it in the manner Wittgenstein describes, since they could, if the 'beetles' are ordinary material objects, refer to their intrinsic properties directly: but let us waive this difficulty in what is after all only supposed to be an illustrative analogy.) If they did not believe that beetles were [intrinsically] similar, they would have no use for a language-game whose logic guaranteed nothing about what beetles were intrinsically like. (At any rate, this would be their attitude if their purpose was as I have specified: if they merely wished to refer to beetles qua contents-of-boxes, things would be different.) It is the article of faith 'Beetles are similar' that makes the game worth playing. I have dealings with my beetle in the belief that it is not, intrinsically, peculiar to me. 'But it does not make sense to say "Beetles are similar" if this is intended to mean anything more than "Beetles are all in boxes."' The answer to this objection is in chapter 1.

This question of the importance of subjective properties, even if they belong only to referents rather than appearing in meanings, arises throughout language, if in a less important way than in the case of sensation terms. The general point is that, as far as the mechanics of language are concerned, the private associations and feelings we have as we speak are irrelevant: they are wheels that turn without anything moving with them. Nevertheless they are not therefore negligible for all purposes. Though quality of experience may be logically irrelevant to communication, it may yet be on account of the experience that we want to communicate at all. Admittedly this point is somewhat vague as expressed here: it is difficult to make it clear. But I think it is of fundamental [149] importance. Wittgenstein touched on it (1953, part 11,

WHAT DO SENSATION TERMS MEAN?

[150] ways – the same two ways. Either the *point* of the term in question is (in part) to pick out a particular subjective type of experience, so that the meaning of the term is (partly) intrinsic on the basis of the term's function; or, even if the term has a more extrinsically based function, the relevant subjective experience is implicated by association. This second kind of semantic involvement is parallel (like the first) to what was said above about initially unsemantic teaching links. However functionally unimportant subjective experience may be, it is only natural for it to be involved in the meaning of sensation terms, simply because it is always there. The primary *point* of having a term may be to pick out items that fulfil a particular publicly specifiable function, but the incidental (from the functional viewpoint) subjective properties of these items tend, as a result of the very dependability of their presence, to be inseparably associated with the function their owner performs.

So there are two weaknesses of the method of enquiry by discovering teaching links. First, we cannot tell for sure that the links we discover are semantically relevant; and second, we cannot tell whether intrinsic factors are also important. For example, is it part of the *meaning* of 'pain' that, when in pain, we cry out? Again, itches are sensations that make one (want to) scratch, for sure; but is this the only criterion? Does 'itch' mean

§vi): 'Suppose someone said: every familiar word, in a book for example, actually carries an atmosphere with it in our minds, a "corona" of lightly indicated uses. – Just as if each figure in a painting were surrounded by delicate shadowy drawings of scenes, as it were in different contexts. – Only let us take this assumption seriously! – Then we see that it is not adequate to explain *intention*. For if it is like this, if the possible uses of a word do float before us in half-shades as we say or hear it – this simply goes for *us*. But we communicate with other people without knowing if they have this experience too.' Of course, many of the associations to which Wittgenstein refers are irrelevant even from the standpoint which I am trying to characterise. But they should not be ruled out of court en bloc on account of a tenet of verificationist methodological dogma. Even if they cannot, as I maintain they can, straightforwardly provide criteria for the use of a concept, they can still crucially underlie its use in the way I have just described.

WHAT DO SENSATION TERMS MEAN?

simply and solely ‘a sensation that makes one want to scratch’? [151] Might it not also matter what itches feel like subjectively, intrinsically? Might ‘itch’, in fact, be a co-extrinsic-cum-quality-intrinsic concept? (And are itches and tickles intrinsically differentiated?)

To answer questions such as these, we need to turn to the second strategy outlined earlier for enquiring into the meaning of sensation terms – the device of imagining changes in the state of nature, alterations of natural law. We can ask questions of this kind: ‘Supposing the sensations that lead us to want to scratch became subjectively quite different, would we still be happy to call them itches, meaning by “itch” what we have always meant?’ or ‘Supposing sensations of the subjective quality that now leads us to want to scratch stopped doing so, would they still be itches?’ On the assumption that these suppositions can be understood (not a trivial assumption), then if the answer to the first question was ‘No’, we would have evidence that there is an intrinsic element (at least) in the meaning of ‘itch’; and if the answer to the second question was ‘No’, we would have evidence, if any were needed, for the involvement of the principal teaching link for itches in the meaning of ‘itch’.⁹⁰ (The relative importance of intrinsic and extrinsic elements, if these turn out to be both, is a more difficult thing to establish.) [152]

Changes of nature: pain

This section of the chapter will be somewhat less sketchy than the others, partly because I believe that the methodological tool it exemplifies is an important one for my central task of defending the importance of subjective experience, and partly because it gives me an opportunity to develop the view I have already given in chapter 1 (pp. 13–14) about the meaning of ‘pain’. ‘Pain’ has long been the subject of stimulating philosophical disagreement, and for that reason, as well as because of its great intrinsic

⁹⁰ [And against the relevance of subjective quality.]

WHAT DO SENSATION TERMS MEAN?

interest, I shall use it as a case study to exemplify the method of imagining changes in the state of nature.

It has seemed true to some that for a sensation to be a pain is an entirely intrinsic matter. Kripke (1972), for example, appears to hold this view: I shall examine what he says below. And Hare (1964) holds that there is one of two senses of 'pain' – he dubs it 'pain₁' – of which the view is true, although in another sense – 'pain₂' – certain extrinsic requirements must also (not instead) be met if a sensation is to count as a pain.

The extreme 'intrinsicist' view is certainly not absurd. It would be perfectly possible for us to have a sensation term in our language whose meaning was purely intrinsic – or indeed a whole apparatus of such terms. But the question is, are any of our actual sensation terms of this kind? Specifically, is 'pain' like this? The [153] arguments that it is not are well known: both those which appeal to the manner in which the concept of pain is learnt by children; and those which claim that, learning techniques aside, talk about pains would lose its point if there were no extrinsic element in the meaning of 'pain', and thus that, since we have the concepts we do because of certain purposes which they serve,⁹¹ and which other possible concepts would not serve, the purely intrinsic account of the meaning of 'pain' must be wrong: there would be no *point* in talking about the quality of our pain experiences in isolation from all other considerations, possible though this might be. Rather than enlarging on these familiar arguments directly, I will allow them to arise below in connection with certain imagined changes in the state of nature: for it is in this context, rather than considered in a vacuum, that it becomes clearest what force they have.

First let me characterise briefly the actual state of affairs, deviations from which I shall imagine. Pains (whether necessarily or contingently) are sensations of a certain characteristic intrinsic quality which are caused by certain characteristic stimuli, and have certain characteristic effects, most importantly that we both

⁹¹ [Functional concepts.]

WHAT DO SENSATION TERMS MEAN?

instinctively and deliberately withdraw from their causes, and try to avoid situations in which we would be certain or likely to experience them, for they are unpleasant, we dislike them, and we wish them to cease when they occur. These characteristic attitudes and reactions are useful [154] to us, because most stimuli which cause us pain are, ultimately at least, harmful. The avoidance reaction to pain allows us to escape the stimulus before its harmful effects go too far.

Imagine first of all that the intrinsic quality of pain suddenly stopped marching in step with pain's characteristic outward trappings. This might happen in one or both of two ways. Either we might continue to experience sensations of the same quality as our present sensations of pain, but find that they now occurred outside their accustomed causal context – they might acquire a new characteristic cause, or simply occur at random, in response to no obvious external events; and they might no longer engender the characteristic instinctive pain reactions. Or else we might continue to respond to harmful stimuli in the same way as before, but find that the sensations which heralded the presence of these stimuli, while remaining subjectively of a kind,⁹² had [155] changed their intrinsic character. And it is not inconceivable that both of these two distortions of nature should occur concurrently.

So far, perhaps, the effort of imagination required is not excessive, though some may jibe at the idea that we might continue to have sensations of exactly the same quality as the sensations we now have when we are in pain, and yet not be disposed to take evasive action. (Remember, though, that in the fictional situation envisaged, no action we know of causes the pain to stop.) But if we reinforce and extend this idea by

⁹² This assumes that our actual pain sensations are subjectively of a kind. Is this a rash assumption? I believe that pains do have a common quality, or at the very least a close-knit family of qualities, but perhaps this is an illusion based on whatever the genuine link is that binds all instances of pain together as pains. We have no objective means, as yet, of settling this point. [Not the same as having no means *tout court*.]

WHAT DO SENSATION TERMS MEAN?

suggesting that we might have sensations of just this quality without even disliking them, we begin to wonder if what we are being asked to envisage is any longer a genuine possibility.⁹³ Is it not part of the essential, intrinsic nature of sensations of this quality that they are unpleasant?

It is with this problem that Hare is largely (and Gardiner partly) concerned in his contribution to the symposium 'Pain and Evil' (1964). Is it or is it not inconceivable that we should have a sensation entirely indistinguishable, subjectively, from a sensation of pain, and yet not find it unpleasant?

The intuitive response is to say that it *is* inconceivable. Surely, we feel, being unpleasant *is* part of the intrinsic quality of the sensation – not, as it were, an optional extra tacked on in addition to this intrinsic quality. And to say that something is unpleasant [156] is to say that we dislike it. And to say that we dislike it is to say that, other things being equal, we try to avoid it. So how can we suppose that we might have a sensation of the same quality as our present sensations of pain, and yet not try to escape from it – let alone not dislike it?

Hare thinks that we can suppose this – even more, he thinks that this situation does in fact sometimes arise. He has more than one argument for this view. First, he claims that there could be a degree of coldness of water which I initially dislike, when I swim in the water, but with continued exposure come not to mind – even to enjoy. (Whether ceasing to mind is the disappearance of an extra sensation, or the elimination of a behavioural reaction, or some third thing, he does not say.) He then maintains that the same could happen for pain, though of course here when we had stopped minding we would no longer call the sensation 'pain', since this word implies,⁹⁴ among other things, dislike. The

⁹³ [Cf. Trigg.]

⁹⁴ Hare does hold, as I have already mentioned, that there is a sense of 'pain' ('pain₁') in which no dislike is implied, but these remarks are about what is still the usual sense of 'pain' (Hare's 'pain₂'), even if Hare is right that these two distinct senses exist.

WHAT DO SENSATION TERMS MEAN?

common reaction to pain, Hare holds, is a contingent, not a logical, matter – it is, in fact, a conceivably dispensable (and sometimes actually dispensed with) ‘subjective feeling of dislike’. [157]

Hare is surely wrong to call dislike a subjective feeling. However much we may doubt that a Rylean analysis of mental predicates is true across the board, surely there are many mental terms for which his account is the most plausible, and surely ‘dislike’ is among them. To call dislike a feeling – especially an *extra* feeling on top of the feeling of pain – is surely to make a clear Rylean category mistake. Dislike is not a sensation at all, but a behavioural disposition.⁹⁵

We are heavily predisposed by long habit to assume that we are bound to dislike pain-like sensations. But we should not reject the conceivability of a counterfactual hypothesis simply out of force of habit. And yet, is it *only* habit in this case? Does the difficulty of imagining undisliked pain-like sensations show merely how constricted our powers of imagination are? Or do we have here a case of natural necessity?

I do not know the answers to these extremely difficult questions. It does seem clear, at least, that there is no purely *logical* bar to there being undisliked pain-like sensations: to this extent Hare is right. It is not self-contradictory, surely, to say ‘I am now experiencing a sensation which is qualitatively indistinguishable from a sensation of pain, and yet I do not dislike it.’ But what is less clear is whether, in cases where we have ordinary pain sensations, and so *do* dislike them – whether in *these* cases the dislike is conceivably separable [158] without the intrinsic quality of the actual sensation in question being affected.⁹⁶ Here reason and intuition part company. Rationally, since dislike is not a

⁹⁵ [Though of course disliked sensations do have characteristic qualities *in virtue of which* we dislike them.]

⁹⁶ [Why should this be any different? Only in that it’s psychologically more difficult to imagine the counterfactual case when we’re actually experiencing the sensation in question.]

WHAT DO SENSATION TERMS MEAN?

sensation, we should allow the possibility of pain-like sensations that are not disliked. But intuitively, the unpleasantness of a pain seems part of its very subjective essence.⁹⁷ I suppose that the correct response to this dilemma is to subject intuition to the schooling of reason: but I am not confident, in this instance, of the pupil's capacity to receive instruction.

Hare's other main argument is that the threshold of dislike is above the threshold of the pain-like sensation. We can induce in ourselves sensations which are recognisably pain-like, yet which are too mild to be disliked. Hare adduces as evidence certain experiments in which people were subjected to varying amounts of pain. Their subjective reports showed that, though the description 'pain' was largely confined to unpleasant experiences, there was a range of milder sensations 'continuous in most ways' with the range of pain sensations, except for not being disliked. C. A. Keele is quoted: 'the element of unpleasantness seems to be superimposed on a sensation which runs through the whole range'. Perhaps, Hare suggests, we might deliberately raise our threshold of dislike further and further up this range, until sensations which originally caused us to recoil could now be received with genuine indifference.

Again, there seems no obvious rational objection [159] to this suggestion, and so perhaps it should be accepted, despite the undoubted strain it puts on the imagination. But in any case it is not absolutely necessary, for our present purposes, to make a final decision on this point, for the following reason. What we are being invited to do, first of all, is to imagine what we would say if we had a pain-like sensation in an uncharacteristic context. Doubtless the maximally uncharacteristic context would be one in which such a sensation (a strong one – the difficulty with a weak one is much less severe) occurred without our disliking it. But fortunately it is open to us to shelve this problem, and imagine that the context of our fictitious sensation is uncharacteristic in other ways than because the sensation is not

⁹⁷ [Imagine others liking what you dislike.]

WHAT DO SENSATION TERMS MEAN?

disliked. We can imagine, for example, that the sensation has no deleterious cause, or no regular type of cause at all; that behaviour which normally alleviates pain has, in this case, no effect, or even makes it worse; that the sensation comes and goes in no particular pattern, without any fixed relation to the cycle of the day, or to what we are doing: such abnormalities as these are quite sufficient, even if dislike remains, for the question to be raised whether, though indeed our sensation would be qualitatively like sensations of pain, we would in fact call it a pain, at any rate without serious reservations.⁹⁸

I should say, to avoid misunderstanding, that these hypothetical abnormalities are supposed to be (hypothetically) real, not apparent. Of course, if in practice we came [160] across such mysterious pain-like sensations, our first response would be to say that they were imaginary or unexplained. We would assume that it was our ignorance that made the sensations seem not to have an ordinary causal context, not that the sensations themselves were genuinely outside the normal pattern of nature. Such is the case today with psychogenic or hysterical pain. But it is part of my hypothesis that the pain-like sensations do actually occur stripped of their accustomed context: they never have an injurious cause; there is nothing, discovered or discoverable, which would count as an ordinary explanation of their occurrence.⁹⁹

What, then, would we say if such unnatural pain-like sensations occurred? There are three obvious sorts of answer to this question, corresponding to the three main types of theory about the meaning of sensation terms that I have already described. The first answer is that these hypothetical sensations are definitely pains: the meaning of 'pain' is intrinsic, so that

⁹⁸ [In fact this is a *clearer* from in which to raise the question of the role of the links.]

⁹⁹ [It's as if the pain centres, while still delivering the same subjective goods, were differently hooked up. The difficulty in all such imaginings is that subjective quality seems to be affected by objective factors: cf. pain vs red.]

WHAT DO SENSATION TERMS MEAN?

whatever happens to the public surroundings, we continue to feel pains just as long as we continue to have experiences of the appropriate subjective quality; and if we do not have experiences of this subjective quality, then whatever the public circumstances, we are simply not in pain. That is one simple response. The other simple response is to say that the hypothetical pain-like sensations are definitely not pains: a pain is just whatever sensation occurs in the right extrinsic [161] context, so that a sudden change in the subjective quality of pain would not affect its continued status as pain, and the transference of pain-like sensations to a different context would ipso facto ensure that they were no longer pains: there would be no question of calling pains sensations of the same subjective quality as sensations which had hitherto been called pains, if these sensations did not occur in the relevant setting. (This is the extreme version of C-subtle, already dealt with rather summarily in chapter 2.) The third, more circumspect reply is to confess that one would not be sure *what* to say in the new situation: it would be true neither that we *were* still in pain, *tout court* and unproblematically, nor that we weren't. In one way we would be, in another not. According to this view, our concept of pain looks two ways, to the intrinsic quality of the sensation, on the one hand, and to the extrinsic context in which it occurs, on the other. If one of these aspects of pain were missing, we would not have an easy way of describing the situation. The concept of pain is tailored to the description of the world as it is, not as it might conceivably become. We might in the imaginary predicament be prepared to speak of having the type of sensation we used to have when we were in pain – but that is all.

There are difficulties for both of the simple responses. The extreme C-subtle type of response, as I mentioned above, has already been discussed in chapter 2. So let us look at the extreme intrinsicist response and see what the [162] difficulties for that are.

I have said that Kripke is among those who adopt the intrinsicist view, and I shall now examine a few remarks he makes towards the end of 'Naming and Necessity' (1972) about pain. I

should say at the outset that it is no part of Kripke's purpose in this article to adjudicate between rival theories as to the meaning of 'pain', or indeed even to mention any theories other than his own (he does in fact mention one alternative). He is concerned, rather, with the problems that arise from trying to maintain an identity between pain and the stimulation of 'C-fibres'. Equally it is no part of my purpose here to evaluate Kripke's objections to a physicalist identity thesis of this kind (though I do happen to believe that they lack force, since, surprisingly in the light of his earlier remarks, he finds it problematic that the identity appears to be – but isn't – contingent; when all it really appears to be – and is – is a posteriori). But though our aims differ, Kripke does reveal by some of the points he makes that he has a definite view about the meaning of pain, a view he assumes for the purposes of his main argument, and that this view is strongly intrinsicist. I shall subject this view to closer scrutiny than it invites (given that it only emerges apropos of other matters): if Kripke were defending it directly, doubtless he would present arguments designed to meet my criticisms. But it illustrates the intrinsicist position effectively, and for this reason I ask to be forgiven for any unfair dealing. [163]

Some of the remarks of Kripke's with which I am concerned are somewhat ambiguous, as will emerge, but there is one whose meaning is clear (p. 340): 'Pain [...] is not picked out by one of its accidental properties; rather it is picked out by the property of being pain itself, by its immediate phenomenological quality. Thus pain, unlike heat, is not only rigidly designated by "pain" but the reference of the designator is determined by an essential property of the referent.' Kripke's equation of pain's 'immediate phenomenological quality' with 'the property of being pain itself' shows that he is an extreme intrinsicist, at least with regard to pain. This is, as I have argued, a coherent position.

First, though, if we take this view to be correct, to say that 'pain is not picked out by one of its accidental properties' is at least seriously misleading. Why is this? On Kripke's general view of naming, the referents of names are 'picked out' in ordinary

WHAT DO SENSATION TERMS MEAN?

discourse, most significantly when the meaning of a name is being taught, by their most evident properties, which may or may not be strictly uniquely identifying,¹⁰⁰ and which may or may not be necessary: often, it seems, it is by their contingent properties that we pick out and recognise the bearers of names. So, for example, we pick out gold by its yellowness, Moses and Aristotle by their famous exploits, water by its appearance and taste, tigers by their stripes, and so forth. 164

How do we pick out pains? When we are teaching or learning the meaning of 'pain', we are bound to use what for Kripke are pain's 'accidental properties' – its having characteristic public stimuli and symptoms – for we have no alternative. In our own case, after we have learnt the use of 'pain', we do indeed use the intrinsic quality of the sensation. Kripke, however, when he talks of 'picking out' and 'fixing the reference', is not talking (mainly) about the use of terms by people who already know their meaning: he is talking about how the meanings of words are passed on from one person to another, from generation to generation. And in the case of pain, this is done by means of what, for Kripke, are its accidental properties.

So the only way to make sense of what Kripke says here is to take him to be talking about the picking out¹⁰¹ of pain by a person who has already learnt the meaning of 'pain'. Let us then understand Kripke's words in this way. He is then saying, quite straightforwardly, that I recognise pain by its subjective quality, which is a necessary property of it, while I recognise heat by the sensation it induces in me, which is not a necessary property of it (heat being essentially just a certain degree of molecular motion, roughly speaking, which might have induced quite a different sort

¹⁰⁰ Often, even if they are not uniquely identifying in general, they do perform this function *in the limited context* in which the referent is being picked out.

¹⁰¹ There is another possible sense of 'pick out', whereby 'x is picked out by y' is equivalent to 'y is the defining property of x'. But this sense cannot be what Kripke has in mind here, or the notion of being picked out by accidental properties (as he says heat is, for example) would be self-contradictory.

WHAT DO SENSATION TERMS MEAN?

of experience without [165] thereby ceasing to be heat). And this is perfectly all right.

It remains true, however, that the meaning of 'pain' is taught and learnt by extrinsic means. This is not of course in itself necessarily an objection to the intrinsicist view of the meaning of 'pain'. As I have said, though it is necessary to admit the role of public phenomena in the use of sensation words, it does not follow from such an admission that such phenomena feature in an analysis of the meaning of these words. In Kripke's phrase, their function might be simply to 'fix the reference' of the words. Nevertheless, it can be argued that these phenomena have a role in the concept of pain (at least), which goes beyond the mere provision of a convenient means of picking out a referent. And it is on this foundation that an argument against a strict intrinsicist account of the nature of pain will have to be built.

What is this extra role? This will come out if we look at one or two other remarks of Kripke's, somewhat less lucid than the one just discussed.

On p. 335 Kripke raises the question whether a particular pain sensation might have existed without being a pain, and answers it in the negative. He asks 'Can any case of essence be more obvious than the fact that *being a pain* is a necessary property of each pain?' What is this supposed to mean? The disanalogy is supposed to be with a case like that of heat, where – to use a parallel locution – being a sensation of heat is not a necessary [166] property of each sensation of heat. But this way of talking is misleading, and the disanalogy does not come out clearly. Of course, at the most trivial level, given that there is a pain, then it is a pain, whatever the true analysis of 'pain' may be. But at this level the same goes for a sensation of heat. So it cannot be at this trivial level that Kripke intends his question to be taken. Rather he must mean something like this: given that we have a sensation of a certain quality, then it is necessary that it be a sensation of pain, since what makes a sensation a sensation of pain is just that it has the appropriate phenomenological character; whereas it does not follow necessarily from the quality of any sensation that

WHAT DO SENSATION TERMS MEAN?

that sensation is a sensation of heat, since what makes a sensation a sensation of heat is that it is caused by the appropriate kind of molecular motion, not that it has a certain subjective quality. This, however, is a great deal more controversial, and begs the question in favour of the intrinsicist view. We might rather, and more plausibly, wish to say that the main reason we pick out a class of pain sensations at all is that they fulfil a crucial function in our lives, warning us of impending damage and enabling us to avoid it. If it were not for this function, we would have less interest, at any rate, in having a special word¹⁰² in our language for that [167] particular subjective type of experience which, as it happens, invariably performs the important function. I say 'less interest' rather than 'no interest' deliberately. It seems hard to deny that, even if sensations qualitatively like those we now have when we are in pain were to occur in a new context, we should still notice them for their peculiar quality alone. Though if Hare's supposition that we might conceivably have the same subjective sort of sensation, and yet not find it unpleasant, not dislike it, were allowed, then there would be still less reason for our noticing pain-like sensations as a distinct group. But having shelved the issue of whether his supposition makes sense, we must allow that, even if the present biological function of pain-like sensations is crucial to their collective identity *as pains*, nevertheless if they were to be severed from this function, they might still retain a perfectly clear claim to a collective identity *as*

¹⁰² The world is full of possible patterns, classes, families, groupings that we don't pick out, or even if [167] we do pick them out, don't label. There is a very natural tendency to suppose that the fact that a particular group of items does have a collective name allotted to it shows that in some mysterious way it has a more solid existence than some other potentially nameable but actually unnamed group, whereas all it may really show is where our interests lie. Perhaps Kripke is influenced by this tendency in his supposition that a sensation of pain has an inalienable identity as a pain. It *may* rather be the case that pains are linked to one another by some common property whose absence in some possible world would leave them looking like a very motley band. (If pain is thought of by Kripke as a natural kind, can we not so specify possible worlds that natural kinds do not necessarily recur there?)

WHAT DO SENSATION TERMS MEAN?

sensations of a distinctive subjective type. But this, of course, falls short of endorsing Kripke's implicit [168] claim that they would still, in virtue of their intrinsic quality, be pains.

So if what makes a sensation a pain sensation is (at least *partly*) this functional role, it will not be true in any non-trivial sense that 'being a pain is a necessary property of each pain', or that 'to be in the same epistemic situation that would obtain if one had a pain *is* to have a pain' (p. 339):¹⁰³ a sensation of the same subjective quality might have occurred in an inappropriate context, in which case it would not have been a pain, or at least not straightforwardly so. That a sensation has a certain subjective quality may certainly be a *necessary* condition of its being a pain – indeed I too hold that this is so – but this may nevertheless not be a *sufficient* condition.

A similar infelicity seems to afflict another remark Kripke makes, on p. 337: "‘pain’ is a rigid designator of the type, or phenomenon, it designates: if something is a pain it is essentially so, and it seems absurd to suppose that pain could have been some phenomenon other than the one it is.' Well, it all depends what he means by 'type' and 'phenomenon'. If the type is intrinsically defined, then it can plausibly be argued, as above, that 'pain' is not a rigid designator of such a type – or at least not merely this. That is to say, it is not true of all possible worlds that a sensation of this intrinsic type will, unproblematically, just be a pain, whatever the surroundings in which it occurs. [169] Again, if 'phenomenon' is taken to refer to a sensation only qua having a certain subjective nature, then of course it is *not* absurd to suppose that 'pain could have been some other phenomenon than the one it is'. If 'pain' meant, for example, 'whatever sensation warns us of injury', then it could have been a different phenomenon just as a sensation of heat could.

If on the other hand all Kripke were saying was 'Whatever the true theory as to the meaning of "pain", it is true that if something is indeed a pain, then it is essentially so, and is rigidly

¹⁰³ [Unless the epistemic situation is so broad as to include teaching-links.]

WHAT DO SENSATION TERMS MEAN?

designated by the term “pain”, then what he is saying, if not totally trivial, boils down to the point that ‘pain’ is a rigid designator – while ‘sensation of heat’ is not. *Even this*, though, is contentious. The reason ‘sensation of heat’ is not rigid is that it does not designate the same subjective type of sensation in all possible worlds: if our nervous system were different, we would have a different sensation of heat. ‘Pain’, Kripke would argue, *is* rigid because it does designate the same subjective type of sensation in all possible worlds.¹⁰⁴ But, once again, this is to beg

¹⁰⁴ Is pain a natural kind (even on Kripke’s view of what pain is)? If not, the application of Kripke’s view of naming to ‘pain’ is obscure to me. Fortunately this problem is not directly relevant to my argument. But let me just pose it. Kripke’s notion of rigid designation by general terms applies, by his own admission (e.g. p. 327, line 8), only to natural kinds; and it depends heavily on their special features. For example, the issue of definition can be sidestepped: one can as it were pass the buck to nature, and say (e.g.) ‘Gold is that substance, *whatever its properties turn out to be*.’ All one need do for a natural kind term is ‘fix its reference’, [170] quite casually. This is why Kripke compares natural kind names to proper names, which (he agrees with Mill) have denotation but no connotation. The necessary properties of natural kinds are not part of the meaning of their names, for their names have no meaning. But we have many terms which do not name natural kinds (certainly e.g. ‘bus’, ‘house’, ‘catastrophe’, ‘philosopher’, ‘nonsense’; perhaps ‘leg’, ‘horn’ ‘glass’ and even ‘pain’), and here things are not so easy. If a class of items exists as a class not so much by nature as because we decide to have a name for things with certain properties, then we (not nature) must decide what their necessary (= definitive, not as in the case of natural kinds) properties are. Moreover, does the rigid/non-rigid designation dichotomy still apply? Should we say that ‘bus’ is a rigid designator, while ‘the type of vehicle, other than taxis, usually used for public transport on the roads of London’ or *somesuch* is a non-rigid designator (because in other circumstances this description might have picked out trams)? If so, then ‘pain’ will be a rigid designator even on my view: it will rigidly designate the members held in common by the set of all sensations of a certain subjective quality and the set of all sensations occurring in a certain public context. An analogous case would be ‘tiger’, which, if defined as ‘tiger-in-the-southern-hemisphere’, would rigidly designate tigers in the southern hemisphere. (I am not sure whether tigers would be part of a natural kind, or no natural kind at all, but this would not matter if the notion of rigidity of designation applied indifferently to all designators.) By the same token, if a purely extrinsicist theory were true of ‘pain’, ‘pain’ would be a non-rigid

WHAT DO SENSATION TERMS MEAN?

[171] the question against an extrinsicist (or semi-extrinsicist) theory of the meaning of pain, according to which if a qualitatively different sensation performed the same function, in some possible world, as our sensation of pain currently performs in our world, then this sensation would (have some claim to) be the sensation of pain; and if the subjective type of sensation we at present experience when we are in pain were to start occurring in different surroundings, then it would no longer (simply and straightforwardly) be the sensation of pain.

At any rate, though Kripke's picture of 'pain' as having a purely intrinsic meaning certainly makes sense, there do seem to be considerations which count against its being the whole truth (though it may indeed, if the semi-intrinsic–semi-extrinsic view is right, be part of it), considerations whose importance is reflected by our reaction to imagining that they did not apply. The discomfort we feel when asked to envisage calling pain-like sensations outside their customary causal context pains is evidence that 'pain' is, partly at least, an extrinsic concept.

Similarly, evidence for an intrinsic element in the meaning of 'pain' arises from imagining un-pain-like sensations in the characteristic causal context of pain. Should we call 'pain', without reservation, sensations which, as far as extrinsic criteria went, were certainly pains, but which had an experiential quality quite other than that to which we are accustomed? It seems to me that the answer to this is 'No.' This is briefly put because [172] there is not really a great deal more one can say about it. It is a matter of intuition. But as I shall be explaining in a little more detail below, this does not mean that the answer to the question should be rejected, either as false or as untrustworthy. The same goes for the question 'Would we say that such a sensation was straightforwardly *not* pain?', to which my answer would also be in the negative. As with the previous imaginary change, we are in a cleft stick if the pattern of nature alters.

designator (like 'occupant of the southern hemisphere'). There is room for much clarification in this area.

WHAT DO SENSATION TERMS MEAN?

There are other possible changes in the state of nature about which it is perhaps not so easy to be definite: it is not so clearly true that we would be the victims, in the way described above, of an irresolvable (save by fiat) dilemma about whether or not we should use the word 'pain' in the counterfactual situation envisaged.

Suppose, for example, that the quality of the sensations we have in the pain context were to change not suddenly but gradually, perhaps almost imperceptibly. This would not be such a baffling predicament to be in. Each day's pain (speaking, for simplicity, in strict C-subtle terms) would be recognisably related to the pain of the day before, so that even by the time the quality of the sensation had changed sufficiently to be substantially different from the quality of the sensation with which the process of change began, the temptation to deny that the sensation was a pain would certainly be less strong [173] than in the case of a sudden change. But would this temptation be entirely absent? Would we say 'The quality of my pains has been gradually changing' or 'I have been finding that the sensations which injuries and suchlike produce in me are becoming less and less painful'?

This question is certainly hard to answer, and the difficulty in answering it might make one sympathetic to Hare's move of saying that 'pain' has two senses; indeed it might persuade one to add a third sense in which the quality of a sensation is not only not the only relevant criterion of whether that sensation is a pain, but no criterion at all. Nevertheless, I think that the dilemma of what to say about a case of this kind is less severe than it may so far appear.

It is a matter of how radical we imagine the change in the quality of the sensation to be. If it is noticeably a change, and yet the new sensation is still obviously related in quality to the old one, we may stick to 'pain' without much hesitation: our concepts are surely not so inflexible as to rule out comparatively minor changes of this kind. We may talk of a requirement that a sensation of a certain kind should have a definite subjective

WHAT DO SENSATION TERMS MEAN?

quality, but this must not be taken too strongly. On a realistic interpretation, the requirement means only that the sensation should fall within a definite subjective range (though even this loosening of the rules would not eliminate the problem of the borderline case). [174]

If, on the other hand, the change is so great as to take the sensation quality unrecognisably far from its origin (though each step is recognisably related to the preceding one), then I think there would be no doubt that we would not be willing to say, without qualification, that the sensation was still a straightforward sensation of pain. For practical convenience, of course (see further on this below), we might talk in the same old way: but this would not obliterate our awareness that things had changed in such a way that by continuing to use ordinary vocabulary we were, strictly, redefining our terms to suit the new situation.

Suppose, lastly, that we are concerned not with a case of change, but with a child who is born with normal pain behaviour but with accompanying sensations of grossly abnormal kind. Some would deny that this supposition made sense – that the notion of a difference in sensation without a difference in behaviour can be understood. But I think we can understand the hypothesis if we imagine that for some reason the central neural networks responsible for responding to injury had a different structure and/or location in this child as compared with a normal child. The input–output connections or correlations would have to be normal, since *ex hypothesi* the child's pain behaviour is normal: the child would respond to painful stimuli with the usual repertoire of distress signals and evasive action. But the route taken within the [175] central nervous system might be different; perhaps, as a result of local cerebral incapacity at the normal site of pain representation, another (natively non-specific) brain area has been called into service, an area with different structural and/or physical properties that yield a different quality of experience in response to injurious stimuli. (A more extreme case of the same general type of discrepancy is provided by the 'visual' system of bats: the whole neural system is radically unlike our

WHAT DO SENSATION TERMS MEAN?

visual system, from the source and nature of the stimuli and the location of receptors onwards, and yet the information the bat acquires is sufficiently analogous to that provided by human vision for the question 'Do bats see?' to be neither absurd nor easy to answer.)

Should we say that this child feels pain? We may allow, certainly, that the difference in question might not show up,¹⁰⁵ for there are no memories of a previous sensation, as there are in the cases of (sudden and gradual) sensation change considered above, against which to compare the abnormal sensations. But that something does not show up does not necessarily mean that it is not there. So [176] the question I am asking comes to this: if, counterfactually, we knew that this child's 'pain' sensations differed in quality from those which normal human beings have in such situations, should we continue to say, quite readily, that he was in pain? (This would be the recommendation of the moderate version of C-subtle, according to which all that is required is that a sensation should, in addition to satisfying the relevant public functional requirements, be of a consistent subjective quality *for a given person*.)

I think we should not. We might decide, of course (as mentioned in connection with the previous example), for practical purposes, to speak of the child in the ordinary way. But this policy of convenience need not alter the fact that, when pressed, we would allow that this abnormal child's 'pain', since it feels different from our own normal pain, is not, strictly speaking, bona fide pain. Equally, if there were a child of whom we knew that he experienced sensations similar in quality to those we

¹⁰⁵ There is more than one way in which it might in fact show up. First, on some views of the nature of the mind, advances in neuro-physiology might enable us to detect such aberrations: this would be so if the abnormality was caused in the way described above. Also, it seems unlikely that a radical difference in the quality of a sensation will not have detectable public repercussions somewhere along the line. In the present case, for example, it would be surprising if there were not *some* detectable difference in the child's attitude to pain, or to a sub-group of pains.

WHAT DO SENSATION TERMS MEAN?

experience when we are in pain, but that these sensations were not caused in the usual way, and produced unusual reactions, then although we might find it best to speak as if he straightforwardly felt no pain when these sensations occurred, we should do so against the background knowledge that, strangely, the child did from time to time, out of context, experience sensations of a pain-like quality. [177]

Admittedly I am relying here, and often above, on my intuitions about what we would say in certain circumstances. The argument may seem breathtakingly brief, or even entirely lacking. Intuitions may seem poor grounds upon which to construct a view of the meaning of sensation terms. But they are, in fact, in the last resort all we have to go on. They underlie both of the methods of discovering the meaning of sensation words that we have been considering. More than one theory is logically compatible with all the observed facts of normal discourse: so we have to abandon empirical evidence and turn to intuition in order to choose which theory to adopt. It would be possible, of course, to raise the statistical significance of the intuitions on which I am relying, by seeing how widely they were shared; and to test them in action, as it were, by trying somehow to discover what people say in actual situations (in so far as the relevant actual situations occur), rather than merely when asked certain abstract, hypothetical questions. So certainly my fieldwork is deficient. But there is not another *kind* of evidence which I am neglecting.

How far have we got in the search for the meaning of 'pain'? We have used imagined changes in the state of nature to pit extrinsic factors against intrinsic factors, and the upshot seems to be that neither a purely intrinsicist view, nor a purely extrinsicist view, will do for 'pain'. By process of elimination, then, if for no other reason, the intermediate position seems preferable, [178] and there is no obvious internal difficulty in adopting it: the meaning of 'pain', that is, is neither wholly intrinsic nor wholly extrinsic, but an indissoluble amalgam of the two. Should there be a radical change of nature such that the concept no longer had a straightforward application, we should be obliged, in the first

WHAT DO SENSATION TERMS MEAN?

instance at any rate, to resort to circumlocution in order to describe our situation accurately. Later, doubtless, a new terminology, purpose built to suit the new situation, could be developed. We might even use the old terms in new senses, for they would no longer be usable in the old ones.

We could go on and probe deeper than this, in search of a more detailed specification of the meaning of ‘pain’, perhaps along the lines sketched in the section entitled ‘A framework of enquiry’ **above**. We could ask, for example, whether a mere decrease in a sensation’s subjective intensity is sufficient to stop it being a pain, or whether all sensations of a pain-like quality, whatever their intensity, are necessarily pains (given, of course, that their objective context is appropriate). This would be an expansion of Hare’s enquiry mentioned above. We could ask whether it is more crucial to a pain that it is caused in a certain way, or that it has certain effects – that is, is ‘pain’, as far as its extrinsic aspect is concerned, fundamentally a causal pre-extrinsic or a causal post-extrinsic concept? Here we might be helped by the examination **[179]** of pain’s teaching links that we undertook in the previous section. We might expect it to be a predominantly preextrinsic concept, given the plausible hypothesis that the basis of pain talk is our need to avoid harmful contact with our environment; though, by the same token, it is important that evasive action should be effective, and this would introduce a post-extrinsic element. If this expectation were fulfilled, the importance of some of pain’s more striking effects or expressions – crying, shouting, wincing etc. – would be pedagogic rather than semantic; their disappearance would not affect the concept of pain as severely as would the other unnatural eventualities we have entertained. The prominent pedagogic role of these effects is not surprising, since especially among children these kinds of pain behaviour are usually more evident, and more constant from case to case, than the often silent and very various stimuli that elicit pain, and the scarcely less, various measures that alleviate it. What might be more surprising to some would be that these types of behaviour were not as intimately involved in the meaning

WHAT DO SENSATION TERMS MEAN?

of 'pain' as is now usually supposed, that their undoubted role in the teaching of the concept was not mirrored in its semantic analysis.

Further detailed enquiry along lines of this kind might employ further consideration of changes in the state of nature, more localised, specific changes than those considered above, changes tailored to the particular detailed questions at issue. What if children reacted [180] to pain differently? But enquiry at this level of subtlety is, as I said at the outset, beyond the scope of this thesis: if I have indicated some directions in which the discussion might proceed, and some methods by which possible answers might be found and assessed, that is enough. I shall not here espouse any detailed hypotheses about the microstructure of sensation concepts. What I do want to do, however, is to extend my broad intuition about 'pain', that it is an extrinsic-cum-intrinsic concept, into a general hypothesis about sensation concepts, a limiting structure to which I believe all individual analyses will be found to conform.

A general hypothesis

At the beginning of this chapter I mentioned that the question 'What do sensation terms mean?' could be taken both in a general and in a particular sense. The present hypothesis represents as much of an answer as I believe to be possible to the question understood in its general sense. It is offered, in true hypothetico-deductive style, in advance of a consideration of more than a fraction of the evidence, but it has the merit, also Popperian, of having sufficient content to be strongly falsifiable, and correspondingly informative if true. It is sufficiently definite to rule out the extremes of pure intrinsicist and thoroughgoing C-subtle, but sufficiently flexible to allow for the variety and fluctuation which [181] surely exist.

First, a recapitulation of the options: the facts about sensations and the potentialities of language being what they are, there are a number of types of sensation terms which, a priori, seem

WHAT DO SENSATION TERMS MEAN?

possible. Sensations could be identified, distinguished, classified and named entirely in terms of their intrinsic, private quality (though of course extrinsic factors would still serve as teaching links, to fix the reference of the terms); entirely by the public contexts in which they occur; or jointly by some combination of the two. In whichever of these three¹⁰⁶ ways the meaning was determined, it would be appropriate to talk of a 'sensation term', since the presence, at least, of a subjective experience would be required, even if its quality was immaterial – and surely, without a subjective experience of *some* kind there can be no sensation.

Such evidence as we have considered, much of it admittedly intuitive, favours the hypothesis that the sensation terms that feature in our actual language have meanings that require, in different proportions from case to case, the satisfaction of both private and public [182] criteria. And by 'private criteria' I mean something more specific than the basic requirement that a sensation be a subjective experience: *that* would hardly be a very daring hypothesis. I mean a requirement that a sensation of a particular kind should have a certain definite *kind* of subjective quality. The exigencies of practical living may in some instances¹⁰⁷ cause the private criteria to recede into the background of our awareness; but the fact remains, it seems, that when we are pressed to admit that these private criteria have faded out of the picture altogether, or were never part of it, we resist this pressure, with some conviction. Their satisfaction is taken for granted rather than being optional or accidental.

Having enunciated this hypothesis, I shall finish this chapter, and the thesis, by showing very briefly how it might apply in one or two cases besides that of 'pain'. This I shall do in the course of

¹⁰⁶ Behaviourism may appear to be a fourth alternative. But 'sensation term' is a misnomer for a behaviouristically analysed term, for the reason given at the end of the sentence.

¹⁰⁷ But only some. When we behave sympathetically towards some-one in pain, for example, we do imagine their sensation; we are not merely trying to stop their pain-behaviour. A similar point might be made about a remark like 'Just look at the colour of that sky!'

WHAT DO SENSATION TERMS MEAN?

making some further remarks about the application of the framework offered above for enquiring into the meanings of sensation terms.

One of the suggestions I made in the earlier discussion of this framework was that the more general a sensation [183] term, the more important extrinsic factors might be in its meaning, and the more specific a term was, the more likely intrinsic factors would be to come to the fore. (And yet, if my general hypothesis is right, in neither case will one kind of factor entirely oust the other.) Let us see what *prima facie* plausibility this suggestion has in a particular case, by considering the range of sensation terms mentioned in this connection before, the range that starts with the extremely general ‘visual impression’, passes (in one direction) through ‘colour sensation’, and ends up with specific colour sensation descriptions such as ‘sensation of red’.

‘Visual impression’ may seem on first consideration to be likely to falsify my general hypothesis by being an entirely extrinsic concept – as ‘pain’ would falsify it if, as some believe, it is an entirely intrinsic concept. Visual impressions may be, for each of us, phenomenologically of a kind,¹⁰⁸ but what, we might

¹⁰⁸ Even this is only true in a weak sense. To put it vaguely (it is hard to do otherwise), visual impressions are experienced *as external scenes*, and not as affections of our eyes, or of any other part of our bodies. To this extent they have no subjective quality, only objective qualities. But this dichotomy is oversimplified, for there does seem to be a sense in which visual impressions have phenomenological properties, which may change for an individual (if, for example, he dons sun-glasses), or differ between individuals (if one of them is, for example, colour-blind). I wish I could be clearer about the nature and implications of this aspect of visual impressions. In a sense, visual impressions are not sensations at all.

After writing this I discover that Grice (1962) has precisely the same difficulty: [184] ‘such experiences (if experiences they be) 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 or 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 what we feel’ (p. 144). Actually, I think, feeling (i.e. by touch)

WHAT DO SENSATION TERMS MEAN?

ask, has this got to do [185] with their being visual impressions? If another person, or another species, received just that information which comes to us through our eyes, by means of a subjectively different range of sensations – say those sensations which, as things are, we call auditory – would we have any hesitation in calling these sensations visual impressions?

As in the case of ‘pain’, three replies are possible: that being a visual impression is simply and solely a matter of inner quality of experience; that a sensation’s identity as a visual impression is entirely determined by extrinsic criteria – whether by the organ through which the sensation comes, or by the external mechanism (light rays) which brings it to us, or by the kind of information it yields (information about shape, colour, brightness) or by some combination of these (i.e. it is a causal pre-extrinsic concept); and, finally, that a visual impression must have both a subjective quality of a certain kind, and an appropriate external cause.

The purely intrinsicist answer would in this case perhaps find few supporters. The exclusively extrinsicist reply might be more popular; but would it be plausible? Imagine that the effect of light on us was such that we discriminated shapes, colours,

is far less problematic than seeing. Indeed we do often misinterpret tactual stimulation as an affection of our bodies.

Grice agrees with me that the difficulty does not make the notion of a special type of subjective experience for each sense senseless: ‘there is a generic resemblance [...] which differentiates visual from non-visual sense-experience. This resemblance can be noticed and labelled, but perhaps not further described. To object that one cannot focus one’s attention, in a given case, on the experience of seeing as distinct from the properties detected is perhaps like complaining that one cannot focus one’s attention on the colour of an object, ignoring its particular colour’ (p. 152).

A related question is this: do we perceive pain as within us rather than as a quality of external objects simply because ‘our pains are on the whole very poor guides to the character of things that hurt us’ (Grice p. 134)? (Contrast smells: ‘Particular types of smells, on the other hand, are in general characteristic of ‘this or that type of object’ (ibid.)’) It seems hard to believe; but it would make sense.

WHAT DO SENSATION TERMS MEAN?

brightnesses and the rest by means of sensations subjectively like those now caused in us by sound waves. For example, the afferent nerves of sight and hearing might be crossed over, so that stimulation of our eyes caused (what we now call) [186] auditory sensations. If full advantage were taken of the potentialities of the new situation, pitch, volume and timbre might substitute for the three dimensions of normal vision, and tone for colour. Would we call these sensations 'visual impressions', meaning thereby just what we mean by the term as we use it at present? It seems to me clear, again on intuitive grounds, that we would not, which suggests in turn that the meaning of 'visual impression' is not at any rate wholly extrinsic.

Nor, of course, is it wholly not extrinsic. Light waves must surely be involved; not to mention eyes. If we received the information that we currently receive by the agency of light, but by bouncing sound waves off the environment, as bats do to find their way around, we would hardly call the resulting impressions visual, even if their subjective quality was as it now is.¹⁰⁹ Nor

¹⁰⁹ But it would be unlikely to be, just as it surely isn't in the case of bats. After writing the above I find that Thomas Nagel (1974) has thought about bats also. He writes: 'we know that [...] bats [...] perceive the external world primarily by sonar, or echolocation, detecting the reflections, from objects within range, of their own rapid, subtly modulated, high-frequency shrieks. Their brains are designed to correlate the outgoing impulses with the subsequent echoes, and the information thus acquired enables bats to make precise discriminations of distance, size, shape, motion and texture comparable to those we make by vision. But bat sonar, though clearly a form of perception, is not similar in its operation to any sense we possess, and there is no reason to suppose that it is subjectively like anything we can experience or imagine' (p. 438). With this I agree. Indeed his whole article is, [187] excellently, very realistic and explicit about the subjective character of experience and the problems it poses. He believes, rightly it seems to me, that it is beyond our ability to conceive what the subjective experience of being a bat engaged in perception is like. And he provides what might be a motto for this whole thesis: 'to deny the reality or logical significance of what we can never describe or understand is the crudest form of cognitive dissonance' (pp. 440–1).

WHAT DO SENSATION TERMS MEAN?

would [187] we continue to speak in the same way if it were our ears instead of our eyes that were sensitive to light.

What are the necessary factors, then, that make a sensation a visual impression? Let us recapitulate the candidates: (i) the introspectible phenomenal quality of the sensation (to the extent that visual impressions have such a quality); (ii) the nature of the information provided (i.e. information about the shape, colour, brightness, position and size of objects in the vicinity); (iii) the external mechanism which provides the information (light rays reflected from objects); (iv) the organ through which the information comes to us (in our case, at least, the eyes), together perhaps with its mode of connection with the brain. It is possible, by more or less strenuous stretchings of the imagination, to suppose of each of these candidates in turn that it is absent [188] in a particular case. For example: (i) the phenomenal quality might be of the kind now called auditory, in the way already described; (ii) the information might tell us rather of sounds, as in cases of synaesthesia (something like this would be the other effect of the imagined neural switch); (iii) we might replace light rays with some complex electronic device which transmitted stimulation direct from the surfaces of objects to the retina; (iv) as Pears (1976) suggests, the eyes of a blind man might be replaced by 'a prosthetic device which worked as hearing aids now work, except that it would have to be more elaborate'.

We must then ask, of each of these cases, whether in the deviant situation imagined we would still call the sensations involved visual impressions. These are difficult questions. Our answer might be different depending on whether the hypothetical abnormality is supposed to be universal, or an isolated deviation from the normal situation with which we are all familiar. And intuitions might vary from person to person. My intuition is that all four factors play their part in the concept of visual impression, although it would be plausible to give factors (ii) and (iii) a greater weighting than factors (i) and (iv), on the grounds that the imagined changes in respect of these two factors produce phenomena least like normal vision. But despite this possible

WHAT DO SENSATION TERMS MEAN?

difference in weighting, it seems to me that none of the four factors is negligible. If one of them were changed for good, the concept of [189] visual impressions might indeed be retained: but it would be a very changed concept. If I am right about this, ‘visual impression’ is an extremely complex notion, not to be easily defined.¹¹⁰

Even if I am right that there is an intrinsic element in the meaning of ‘visual impression’, it must be allowed, as I have suggested, that it has a secondary importance. To the question ‘What makes a sensation a visual impression?’ one’s first response is hardly to fumble for a description of an elusive subjective quality, but to mention (e.g.) light and eyes. This explains why we suspected at the outset that an exclusively extrinsicist account of the meaning of ‘visual impression’ might be possible: the extrinsic determinants, though not the [190] only ones, are predominant. Some, like Pears,¹¹¹ disagree with my intuitions, and feel that the subjective quality of visual impressions is sufficiently unimportant for us to be able to contemplate using the term ‘visual impression’ even if this quality were to change (though

¹¹⁰ Pears (1976) writes: ‘The essence of seeing seems to be mainly functional: it must yield information about the disposition, shape and size of local objects. If we add a stipulation about mechanism, that this information must be carried by patterns of light-rays, we get a plausible account of the essence of seeing. If some creature has a sense which gives it this information in this way, it follows that it is sight. If it has a sense which gives it this information in a different way or gives it different information in this way, it follows that it is not sight.’ But this is too simple. Do eyes and quality of experience really have no part to play?

¹¹¹ He writes (*ibid.*): ‘Visual experiences evidently do not have to possess the phenomenal characteristics of ours. The multifaceted eyes of insects must produce visual experiences of a very different kind.’ But: (a) Why should multifacetedness alter the phenomenal character of visual experiences (admittedly it would alter their *structural* properties, but this is something different)? (b) Do insects have experiences *at all*, with such attenuated nervous systems (compared to ours)? (c) Even if they do, is not the proper standard for visual experience a human one, with the concept being borrowed, in so far as it fits, for non-human cases?

WHAT DO SENSATION TERMS MEAN?

aesthetic¹¹² considerations alone seem to me sufficient to rule this out). At any rate, does this predominance of extrinsic considerations, however total, alter when we move to the other end of this range of sensation terms, say to ‘sensation of red’?

It is well known that ‘sensation of red’ is ‘definable’, nowadays at any rate, in terms of the range of wavelengths of light which characteristically give rise to the sensation of red. Before we knew anything about wavelengths, it could be ‘defined’ in an analogous way in terms of typically red objects. This is not surprising: we must use public devices to fix the reference [191] and discuss the meaning of sensation vocabulary. On the other hand, everyone who is not colour blind will agree that a sensation of red has, subjectively, a very characteristic quality, easily recognised intrinsically, such that if typically red things, or light of wavelength 620 millimicrons, began to give us a colour experience of quite a different subjective quality, we should not be happy to continue to describe these objects, this light, as red. Indeed, one might say that it is most importantly the fact that they cause a certain definite subjectively characteristic experience in us (the same for everyone) that makes red objects red;¹¹³ that it is for the same reason, principally, that light of wavelength 620 millimicrons is called red light. In general, colour is a purely sensibly determinable property. Of course, it is true that red is also essentially a *colour*, so that sensations of red must be caused in the right way for colours: if (all) our red-like sensations were caused by random neural activity in the brain, or by sound waves, or by electrode stimulation, or somesuch, then ‘sensation of red’ [192] (if it were used to describe these experiences) would not be the concept it now is. But this need not affect the fact that, among the essential properties of a sensation of red, the quality of

¹¹² [Non-functional!]

¹¹³ Is a sensation of red whatever sensation is caused by red? Is red whatever causes a sensation of red? Is there a compromise between these two alternatives? Whichever of ‘red’ and ‘sensation of red’ is a rigid designation in Kripke’s sense), what is it a rigid designation *of*?

WHAT DO SENSATION TERMS MEAN?

experience is primary, and the public phenomena in terms of which it can be 'defined' secondary. Again, maybe there is some evolutionary reason why we have a special subjective experience evoked by the range of wavelengths in question, and do not have one of the numerous possible alternative mappings of colour experience on to wavelength: in which case there may be an extrinsic *explanation* of why we (are in a position to) have the concept 'sensation of red'. But this too does not affect the contention that 'sensation of red', as far as its *meaning* is concerned, hooks on primarily to a particular subjective type of experience. Some uses of 'red' ('Red is my favourite colour') reflect this intrinsic criterion more than others ('Connect the red wire to the live terminal'): but even if in a particular instance its importance seems negligible, it continues crucially to underlie the totality of our talk about red.

'Colour sensation' falls, as one might expect, somewhere between 'visual impression' and 'sensation of red'. The extrinsic and intrinsic conditions that a sensation has to meet in order to count as a sensation of colour are more equally balanced in importance than in the other two cases. This may be partly because 'colour sensation' is a more general term than 'sensation of red', and so less crucially associated with the range of subjective [193] experiences it covers; and a more specific term than 'visual impression', and so more likely to be associated with the narrower range of sensation qualities to which it (indifferently) refers. But more importantly, its intrinsic content is less crucial because we can imagine someone with abnormal vision who is sensitive, say, to ultraviolet light, and has a special colour sensation to match – just as the colour blind person can imagine (and what is more he knows it to be true) that people with normal colour vision have colour experiences in their repertoire that are missing in his. So although it may be more important to a colour sensation than to a visual impression what its subjective quality is – as I have said, the range of options is smaller, and this in itself brings intrinsic considerations more into the limelight, if only by association – its subjective quality is not

WHAT DO SENSATION TERMS MEAN?

as specifically constrained as is the quality of a sensation of red. So 'colour sensation' seems to be both a less intrinsic concept than 'sensation of red', and a more intrinsic concept than 'visual impression', and thus to fall somewhere in the centre of the range of terms we are looking at.

Why should it be that intrinsic considerations become progressively more important as we move from 'visual impression' to 'sensation of red'? At least one contributory cause, presumably, as I have just mentioned in connection with sensations of colour, is the very [194] generality of the one term, and the very specificity of the other.¹¹⁴ A general term covers a wide range of individually different items, and for this reason alone it is unlikely to be an important part of its meaning – and so of the essence of the class it comprehends – that a definite type of subjective quality is involved. There may indeed be – I believe there are – general subjective limitations which all (normal) visual impressions essentially observe, but there is no narrow requirement that they have a particular definite subjective quality. It may be said that this shows only that the required intrinsic quality is generic – disjunctive – not that it is unimportant. But, to repeat, the generality in itself creates a kind of latitude which weakens the subjective requirement somewhat. Moreover, the disjunction is open-ended. If I want to marry a particular girl, and only her, my personal requirement is totally exclusive; if I want to marry a girl, but any old girl will do, then although my determination not to marry a man may be as fixed as my determination not to marry anyone other than the one special girl in the other case, naturally I will be said to be less fussy in this second case. [195] For example, I may end up marrying someone whom, at the time I draw up my marital conditions, I have never

¹¹⁴ This is not as trivial an observation as it may sound. As will emerge below, increasing specificity need not necessarily mean an increasing emphasis on intrinsic criteria: it might rather be the extrinsic criteria which became more specific, while the intrinsic criteria stayed the same. But, as I go on to say in the text, it is more likely that sensation terms will be intrinsic in proportion to their specificity.

WHAT DO SENSATION TERMS MEAN?

met. So it is with visual impressions: we can imagine that an unexpected visual impression might turn up – but an unexpected sensation of red (unexpected qua sensation of red that is)?

Another factor contributing to the relative unimportance of intrinsic elements in the meaning of ‘visual impression’ as compared with ‘red’ may be that the primary principle according to which visual impressions are grouped together is not a subjective one. Just as (if I am right) a sensation of red is a certain type of subjective experience first, and caused in a certain way second, so a visual impression is caused and/or mediated in a certain way first,¹¹⁵ and has one of a large range of subjective qualities second. In both cases the secondary factors are semantically implicated; but they still remain secondary.

On both these counts sensations of red (and, to a lesser degree, sensations of colour) are different from visual impressions. All reds have a specific quality in common, and it is in virtue of this quality, primarily, that they are grouped together, even if, as I have said, science has discovered an extrinsic criterion to double [196] for – even in some walks of life to supplant – the intrinsic one.

There may be further considerations affecting the relative importance of intrinsic criteria at different points on the range visual impression/sensation of colour/sensation of red. Also, even if I am right about this range, it may be a special case – other ranges of sensation terms may not betray such a neat pattern, though prima facie it seems likely that many of them will, since such a rationalisation of the pattern as I have offered would apply equally well elsewhere. But whichever way it turned out, would it have any obvious significance except in terms of pure compilation?

Another suggestion I made in introducing my ‘framework of enquiry’ was that it might be possible to arrive at some generalisation about what is involved in moving from one specific

¹¹⁵ These may come to the same thing, since we may group causes of impressions together on the basis of their sharing a mediator.

WHAT DO SENSATION TERMS MEAN?

sensation to another within a certain genus. This suggestion, however, seems on closer examination to be less fruitful than the previous one: it comes up against the diversity of sensations. Sometimes the quality of the sensation, as well as its cause, needs to change before its name changes, as with sensations of colour. Sometimes (if Hare is right, pain is an example) a mere change of intensity will suffice. In other cases quality (and intensity) may be held constant while variation sufficient to entail a change of name occurs elsewhere. This seems to be, the case with what we might [197] call emotional sensations – sensations of fear, of anger, of hate and so forth. There is considerable experimental evidence (as well as the experiential evidence of those with a reasonably developed capacity for introspection) that the physiological states, and (so) the sensations,¹¹⁶ associated with the various emotions are broadly similar, or even identical, so that what makes the difference is the context – most importantly the context of thoughts – within which the sensations occur. This would be an exception to the observations above about general versus specific sensation terms: that certain definite types of subjective experience were occurring would be as important – even perhaps more important – for emotions in general as for particular emotions, which would be distinguished from one another in other ways, ways extrinsic to the experiences involved.

It is no surprise if no easy generalisation is available about which factors are crucial in transforming one species of a genus of sensations into another. We must surely expect a certain irreducible variety in our sensation terms, because of the variety of the purposes which they, and the sensations they refer to, serve. Even this variety, though, I predict, will be found to occur within the limits of the general hypothesis about the meaning of sensation terms that I have framed. That is, every type of sensation that is singled out by being assigned a name will have some necessary intrinsic [198] characteristic, beyond merely being a *sensation*, either because of our intentions in picking it out, or by

¹¹⁶ [They can be described also in non-emotional terms.]

WHAT DO SENSATION TERMS MEAN?

the process of association described above; and the fact that language is a predominantly practical instrument will ensure that, on the other hand, extrinsic characteristics will always be required of a sensation too. What I doubt is whether a detailed investigation of a large number of sensation terms, from a wide range, will yield anything further beyond a mere detailed list. But even if it does not, I shall be content if my general defence of subjective experience has been successful.

CONCLUSION

What are the important things that emerge from this discussion of subjective experience? They seem to me to be two. First, that interpersonal¹¹⁷ discussion of subjective experience, though extremely difficult if the attempt is made to go into any detail, is nevertheless intelligible as an activity. This is important especially for the following reason: it is a central human need to feel that one can, to some degree at any rate, communicate the subjective contents of one's mind to others.

That one can communicate one's thoughts is not in dispute. But feelings are more difficult to get across. There are many strategies for doing so. Minute description of the circumstances and thoughts associated with feelings is perhaps the best strategy so far discovered: it is constantly in use by novelists as well as in private conversation. One day perhaps, if the identity theory is right, advances in cerebral neurophysiology will give us a new, more direct, more detailed and sophisticated method of communicating the quality of our experiences [200] to one another.¹¹⁸ But methods aside, it might almost be [202] said that

¹¹⁷ That one can record or brood on one's own subjective experiences without involving other people is also important, but perhaps less fundamentally so.

¹¹⁸ Wittgenstein believed, according to Pears, that there is no way to make the description of a sensation which underlies certain teaching links more specific than 'x', where 'x' is a sensation term defined in C-subtle style, other than by reference to our attitudes to people who have the sensation. But if we accept the identity theory of mind in respect of sensations, and allow a link between occurrences of sensations and processes in the brain, we would be in a position, theoretically, to make some progress here. We would be able to describe our subjective experience far more precisely than current methods

CONCLUSION

allow, for neurophysiological terminology has a far greater potential for detail than, for example, the rather gross language we have for describing the behaviour of people who are in pain. Given that one has a stomach ache, and that it is a stabbing or throbbing sensation, one's powers of description are virtually exhausted. But if we read off what is happening in the brain, our powers of discrimination, diagnosis and treatment, to name but three, will be immeasurably improved. Also, if people could be told what states their brains were in when they had a sensation, they could surely learn to recognise a far greater variety of sensations, and far subtler differences between successive instances of similar sensations, because there would be a reliable way of making sure that their use of terms was precise and consistent: whereas the present means at our disposal for checking people's use of sensation terms suffice only to enable them to perform rather broad differentiations. Let me explain this point in a little more detail.

It is well known to learning theorists that powers of discrimination are dependent on contingencies of reinforcement as well as on the structure of the discriminating organism. That is to say, an animal will only learn to use the powers of discrimination it needs, not all those powers which it is physically possible for it to develop. We humans are no exception: the complexity of our sensory organs provides for subtlety in our powers of discrimination far greater than we usually have any need for. We are born with the capacity to manifest many skills we never learn. Skills take practice to develop, and so untapped potential will not feature in our behavioural repertoire. After a certain age, many possibilities which were once open to us become closed: previously available brain space is used up in subserving other pursuits, and cerebral pathways rigidify or atrophy. The failure to develop discriminatory potential is usually due to our [201] lack of a need for its fruits: we do not need to distinguish, in our lives, the multitudinous shades of green which a jungle-dweller can tell apart – and his livelihood or even his life may depend on his ability to do so. Had we been born in the jungle, no doubt we would have developed our powers of discrimination in this direction too. Had we been born in the arctic circle, we would have learnt to tell apart all the many different sorts of snow. It all depends on contingencies of reinforcement: where there is a need for a certain distinction to be made, because the result of making it rightly will be significantly different from the result of making it wrongly or not at all, then, if the system at our disposal permits, we will learn to make the distinction. Differences which begin by being hardly noticeable will gradually strike us more and more forcibly as we are 'rewarded' for noticing them correctly and 'punished' for making mistakes. The physiological basis of such learning is imperfectly understood, but it is a plausible enough phenomenon, and part of everyone's experience: first impressions – of people, places, races – blur distinctions which further acquaintance renders obvious.

CONCLUSION

it is the very intelligibility of communication about private experience that matters more than the availability of any really practical means of engaging in it. The feeling of isolation which results from being persuaded that communication of this sort is not only hard, but senseless – that the very attempt is delusory – is scarcely less terrible than that engendered by belief in solipsism, or in systematic delusion by a malevolent influence, or in the hypothesis that all people besides oneself are sensationless machines. It is important to personal security to feel that others experience the world similarly to oneself – not just that they react visibly in similar ways, that they say the same sorts of things about it, that their behaviour is in many ways broadly predictable; but that the subjective quality of their stream of consciousness is mappable on to their external circumstances and their bodily condition in a way by and large comparable to the way in which one's own experience is so mappable. This is 'community of experience', and belief in it is surely basic in everybody's perception of his social environment.

In the case of discrimination amongst our subjective experiences the reason why we are clumsy is not that better powers of discrimination would be of no interest or use to us. Rather it is that currently available teaching links are far too crude to be used as reinforcers of initially precarious discriminations. If our discriminations are to become more precise, we need to be able, while we are learning, to consult some independent authority on the accuracy of our performance. If we can be told for a large number of trials which decisions are right and which wrong, then we will gradually grow more reliable, until we can eventually dispense with supervision. But if we can never be sure how we are doing, we will continue to do badly. And the latter predicament is the one we are in as far as sensations go: the same links cover a variety of potentially discriminable subjective experiences which we cannot, therefore, learn to tell reliably apart. (Memory is not sufficiently precise to make up for the deficiency of teaching links.) What I am suggesting is that if, while we were trying to improve our powers of discrimination in this regard, a cerebral neurophysiologist was telling us whether or not we were consistently matching sensation descriptions with brain states, we could progress by leaps and bounds. This entirely new system of reinforcement would train us to be far more useful, because more precise, symptom-reporters, besides adding to the intrinsic interest of the inner life.

CONCLUSION

The other important thing that emerges is related to this, and is what this thesis has more narrowly been about. It is that the language in which we talk about sensations – bodily and otherwise – is so constructed that an assumption of what I have just called ‘community of experience’ is built in. The meanings of our sensation terms are such that one at least of the necessary conditions [203] which determine the identity of a sensation is its subjective quality. This is important for reasons similar to those mentioned in respect of the first point. When we feel sorry for someone in pain, we do so because we believe that he is feeling the same way we feel when we are in pain, and we know what it is like. To rationalise our sympathy in this way would be unjustified if it was not the case that to be in pain one needed to be having a certain definite kind of subjective experience. Again, we may ground our moral judgements on the belief that one should promote pleasure or reduce pain: and this basis for morality would be far less intuitively appealing if we did not believe that all men’s pleasure and pain had a subjective overlap. Why should one be concerned to reduce pain behaviour if one did not believe that one was thereby reducing the pain one knows from personal experience?

In short, what I have been presenting as our intuitive beliefs about our subjective experience, and our intuitive model of how we talk about it, seem to me to be exceedingly fundamental in our conceptual scheme. This does not show that these beliefs are true, any more than the deep conceptual entanglements of our belief in free will show that that belief is true (or even intelligible). But it does mean that peculiarly strong reasons will have to be advanced before we will be prepared to abandon such beliefs. And it has been the burden of my thesis that [204] the strongest reasons which have been advanced so far have not been strong enough. [205]

BIBLIOGRAPHY

- Ayer, A. J. (1954) 'Can There be a Private Language?', *Proceedings of the Aristotelian Society*, supplementary volume 28; repr. in Pitcher (1972)
- Berlin, I. (1950) 'Empirical Propositions and Hypothetical Statements', *Mind* 59
- Butler, R. J. (1962) (ed.) *Analytical Philosophy* (Oxford: Blackwell)
- Davidson, D., and Harman, G. (1972) (eds) *Semantics of Natural Language*
- Gardiner, P. L. (1964) 'Pain and Evil', *Proceedings of the Aristotelian Society*, supplementary volume 38
- Grice, H. P. (1962) 'Some Remarks About the Senses', in R. J. Butler (1962)
- Hanfling, O. (1975) 'Privacy', unpublished draft of Open University course unit
- Hardy, H. R. D. (1971) 'The Private Language Argument', unpublished paper
- Hardy, H. R. D. (1974) 'Subjective Experiences', BPhil thesis, University of Oxford
- Hare, R. M. (1964) 'Pain and Evil', *Proceedings of the Aristotelian Society*, supplementary volume 38
- Jones, O. R. (1971) (ed.) *The Private Language Argument* (London: Macmillan)
- Kenny, A. J. P. (1971) 'The Verification Principle and the Private Language Argument (ii)', in Jones (1971)
- Kovesi, S. (1967) *Moral Notions* (London: Routledge and Kegan Paul)
- Kripke, S. A. (1972) 'Naming and Necessity', in Davidson and Harman (1972)

BIBLIOGRAPHY

- Laing, R. D. (1971) *The Politics of Experience* (Harmondsworth: Penguin) [206]
- Legge, D. R. (1975) *An Introduction to Psychological Science* (London: Methuen)
- Locke, D. (1968) *Myself and Others* (Oxford: Clarendon Press)
- Locke, J. (1690) *An Essay Concerning Human Understanding*
- Mill, J. S. (1889) *An Examination of Sir William Hamilton's Philosophy*
- Nagel, T. (1974) 'What Is It Like to Be a Bat?', *Philosophical Review* 83
- Parfit, D. (1971) 'Personal Identity', *Philosophical Review* 80
- Pears, D. F. (1971) *Wittgenstein* (London: Fontana)
- Pears, D. F. (1976) 'The Causal Conditions of Perception', unpublished paper
- Pitcher, G. (1972) (ed.) *Wittgenstein* (London: Macmillan)
- Putnam, H. (1973) 'Meaning and Reference', *Journal of Philosophy* 70
- Ryle, G. (1949) *The Concept of Mind* (London: Hutchinson)
- Strawson, P. F. (1959) *Individuals* (London: Methuen)
- Wittgenstein, L. (1953) *Philosophical Investigations* (Oxford: Blackwell)