Subjective Experiences

Henry Hardy

This is the text of the thesis submitted by the author to the Faculty of Literae Humaniores in the University of Oxford in 1974 in partial fulfilment of the requirements for the degree of BPhil. 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. 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.

A heavily revised version of this thesis, **'Subjective Experience'**, was submitted in 1976 for the degree of DPhil. Each thesis contains material absent from the other, making neither dispensable for a full understanding of the author's position.

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

First Posted in the Henry Hardy Virtual Library 16 December 2018
Reformatted 24 January 2121

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

Contents

Figures

Preface

I should like to thank Stuart Hampshire, Rosalind Hursthouse, Ralph Walker, Mrs D. M. Waller, Lesley Whale and Witney Typing and Duplicating Services, all of whom have influenced this thesis, either directly as supervisors, or indirectly by being helpful to its author in various ways, or both.

H.R.D.H.
Wolfson College, Oxford
March 1974

INTRODUCTION

In what, follows I argue against a certain philosophical theory about the workings of that department of our language which is used to talk about our subjective experiences; and in favour of a more intuitive view. The effect of the philosophical theory to which I am opposed is to deny that a certain intuitive model of the way we talk about our subjective experiences makes sense. According to this intuitive model, subjective experiences are private to the subject: only he is directly aware of their occurrence and of their phenomenal nature. There is no possibility of one subject of experiences comparing his experiences with those of another subject to discover whether or not the experiences of the two subjects are qualitatively the same or similar in the same or similar circumstances. There is no way for a child who is learning language to reveal 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 items being labelled. And yet children do successfully learn the language of subjective experiences; and we do believe that other people's experiences are qualitatively similar to our own in similar circumstances, not to mention that the question whether or not people's experiences are qualitatively similar, and the question whether or not an adult has inferred the presence of the appropriate experience in the child, are questions which make sense, and have answers, whether or not we are in a position to discover them. **[2]**

I believe that this intuitive view is true, and that the need to defend it arises only because of the existence of a philosophical theory which is inconsistent with it. This theory comes not from a

consideration of the way we actually operate when talking about subjective experiences, and of the intuitions we have about such talk, but is required by a certain view of meaning, a view of a verificationist kind, which is false as far as this department of language is concerned. The theory 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, about the inferring of the occurrence of subjective experiences from the observation of objective phenomena. So, according to the theory, 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 be subjective only 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 around to be 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.

The picture painted by this theory may avoid certain problems. But that is of no profit if it is also false. **[3]**

Although I talk of a single theory, there are in fact several different theories, or several different versions of one theory, which espouse views of the kind I have sketched. above. The differences between these theories or versions are less important for my purposes than their similarities, for I hold that they are all wrong, and all wrong for the same sort of reasons. I shall have occasion from time to time to distinguish between some of them, in order to show that certain reformulations designed to sidestep counterarguments are ineffectual. But since I am not directly concerned to give an accurate characterisation of the different forms the theory can take, I shall normally lump them together and talk of them as one.

One proponent of a version of the theory with which I am concerned is Wittgenstein. Pears (1971) discusses Wittgenstein's views on this matter briefly and clearly in Chapter 8 of his book on Wittgenstein. The chapter is entitled 'Sensations'. This chapter is my starting point, and I will build my arguments around it. This explains why, despite my overall title, *Subjective Experiences*; much of my discussion is couched in terms of sensations: but what can be said about, sensations can, *mutatis mutandis*, be said of other subjective experiences too, if there are any (I shall explain this reservation further later).

I begin with 'C', a theory similar to the intuitive view I have mentioned, to which the theory which I am concerned to refute is opposed. 'C' is Pears's name for a theory which, in the eyes of its opponents, 'treats sensations as if they were more like material objects than in fact they are'. According to these opponents, C has as a consequence that the language in which we talk of sensations could not have been, and so was not, taught. **[4]** But this consequence is false: for we do talk of sensations. Therefore either C itself is false, or else the opponents of C are wrong in holding that this consequence flows from it. I take the second view, and wish to adopt a version of C.

Chapter 1

I shall begin by giving my reasons for not believing that C has the consequence its opponents say it has. This means treading again on the well-trodden area of the private language argument, and I shall not dwell on this part of my case in great detail. But my reasons must first be briefly stated, to show that my rejection of the claim that an unacceptable consequence follows from C is not simply a dogmatic assumption. For without the rejection of this claim i would not be able, having argued against alternatives to C, to fall back on a version of C as the true theory. Nevertheless, I shall devote more space to arguing against alternatives.

Chapter 2

Having said why no unacceptable consequence follow from C, I shall then examine 'C-subtle', a leading alternative theory, and show how it is inadequate.

Chapter 3

From this examination will flow a characterisation of just which version of C I wish to adopt. Since it makes a considerable difference which version is adopted (it will determine whether C is true or false), I will occupy a certain amount of space in attempting to specify clearly the important features of my preferred version.

Chapter 4

Next I shall clarify the logic of 'teaching links', as Pears calls them, which play a vital role in my version of C. Mistaken views of their logic, in particular the 'criteriological' view, lead back to a wrong version of C, or, worse, to a form of behaviourism. **[5]**

Chapter 5

Finally I shall look at certain interesting and important possibilities which follow from the acceptance of C, possibilities which, without C, would disappear. It may be a kind of support for C, as against other possible theories of the language of subjective experiences, to specify what these possibilities are: for we may be unable or unwilling to relinquish them. **[6]**

1 THE POSSIBILITY OF C

First, we need a characterisation of C. This I borrow from Pears, p. 150:

Wittgenstein specifies C as the theory which says that a child under instruction can establish the meaning of the word 'pain' for himself by turning his attention inwards on to the right kind of sensation, and affixing the word to it. Here the idea behind the theory is that what the child does is exactly like what he does when he establishes the meaning of the word 'rose' for himself by turning his attention outwards on to the right kind of flower, and affixing the word to it. It is, of course, the teacher who tells him when he is in fact having the right kind of sensation, or is looking at the right kind of flower. Secondly, Wittgenstein specifies C as the theory which says that a person who has a sensation, such as a pain, may know that he has it. This, of course, suggests that there is also the possibility that he might not be sure that he had it, in spite of knowing the meaning of the word 'pain'. Here the idea behind the theory is that the phrase 'I know that' may be prefixed to the proposition 'I am in pain' in exactly the same way that it may be prefixed to the proposition 'This is a rose.' It can be prefixed to the latter proposition precisely because there really is also the possibility that he might not be sure that it was one, and might be mistaken, although he knew the criteria for roses, a possibility which would be realised if, for example, the light were bad.[1]

[1] This might not be possible, for the assumption that material object language is basic, and so potentially self-sufficient, is not at any rate self-evidently true. It might be the case that an ability to use the language of material objects entails an ability to use sensation language. This possibility is further discussed in Chapter 5.

This, no doubt, is an unsympathetic characterisation of C. But it may serve, with various reservations that I shall mention.

The first reservation is that the assimilation of sensations to material objects is exaggerated and oversimplified by this account. 'Exactly like' and 'in exactly the same' are excessive phrases. For we can believe that the language of sensations and the language of material objects have many features in common, enough to make us adherents of C rather than of some theory quite incompatible with C, without supposing either that the two departments of language are entirely alike, without supposing (what *is* supposed by the above account of C) that the **[7]** use of terms for sensations is taught in precisely the same way as, or even in a broadly similar way to, the way in which the use of terms for material objects is taught. We may concede that sensation language is learnt in an indirect fashion, via what Pears calls its 'teaching links' with material object language, without committing ourselves to the conclusion that the logic of sensations is therefore in every way distinct from the logic of material objects. (Even material object language may not always be taught quite so directly as is imagined by this version of C.) Exactly how we may have the concession without the conclusion will become clearer in due course.[2]

Without a reservation of this sort, the suggestion is countenanced that, because sensation language is like material object language, it might exist outside its present setting in the rest of our language, rather as the material object department of our language (supposedly) might. But this is a suggestion that we do not want to countenance, for it opens the door to a kind of argument which our actual sensation language simply doesn't have to meet. This is the argument that an isolated sensation language, conceived according to C, could not possibly be taught, and hence would not be a language at all.[3] Such an argument, even if successful, is beside the point. For it does not count against a

[2] [Does it?]

[3] [Is a metalanguage allowed for teaching? If not, the argument seems clearly effective.]

certain model of our sensation language (say C) that sensation language would not, on this model, be a possibility in isolation from other departments of language: not, that is, unless the model in question entails that even *in* its setting **[8]** the language for sensations is not taught with any reference to its linguistic environment. And it is just this entailment that I am now wishing to excise from C. For our sensation talk is *not* isolated, either in its day-to-day use, or in the way it is taught and learnt; and any effective arguments against any theory of sensation language will have to look at it in its actual context, where it has various circuitous devices at its disposal for getting taught and being monitored in use, devices which would not be available to it if it were obliged to operate on the strength of its own resources.

Pears later calls the version of C I have quoted 'C-crude' to distinguish it from 'C-subtle', which he treats (as can be seen from his labelling) as a variant of C: but I am treating C-crude and C-subtle as opposed theories. It is C-crude towards which I lean, though I do want to modify it in various ways, one of which I have just specified: others, not being germane to my present purpose, I shall discuss in Chapter 3. I relegate to Chapter 2 discussion of C-subtle, since the theory I wish to defend does not, in the respects in which it differs from C-crude, resemble C-subtle: indeed, as I say, C-subtle is the theory against which I am arguing. So that for the purposes of saying why I don't believe that C has its alleged fatal consequence, I shall stick to something like C-crude, with the one important modification already mentioned.

A number of points are made against a theory which, like C, makes the inner referents of sensation terms dominant in their meaning. None of these points seems to me to stick. I mention two main ones.

(i) It is argued that, if C is true, a person cannot learn **[9]** what a sensation term refers to. Let us talk of a sensation called '*x*', if only to avoid talking of pain. *X* is supposed to be a definite sort of thing which a man who has *x* himself can identify. But how could he have learnt which sensation is properly to be called '*x*'? He couldn't have been taught by someone else to identify *x*s, because

no one could know when he had *x* and hence tell him when it was correct to say he had *x*.[4]

Against this kind of argument the assumption that we in fact make (see fig. 1, p. **9**), namely that like sensations go with like stimuli, behaviour or other teaching links, is a quite reasonable defence by any realistic standard (though not by the unnecessarily harsh standard of the verificationist): and just as well, for if this assumption were false, our sensation language (as conceived according to my view) would indeed be, if not impossible to teach, at least useless in one of its most important functions. For one of the things I believe we can do with sensation language is to tell people what we feel, in the sense given to this activity by C: and unless this kind of communication is in fact possible, this sort of use of sensation language is a parody. **[10]**

[11] But the assumption is not false. I shall not argue this point at any length here, but just mention that I accept the 'argument from analogy'[5] as it is called. Wittgenstein asks how we can generalise from the one case (our own) so irresponsibly.[6] The answer is that it cannot be irresponsible to do one's best. There *are* cases where better than analogical evidence is available, perhaps: in these cases let such evidence be made use of. But where other

[4] This difficulty (if it is one) is not confined to sensation language. In the straightforward ostension situations which are supposed to be paradigmatically appropriate for the learning of material object language, there is the problem of ensuring that the pupil interprets the ostension in the same way as the teacher: any act of ostension is in principle many ways ambiguous. This problem may be readily superable by simple objective devices, in which case it would be set apart from the problem which obtains in the case of sensation language: but the problems are, nevertheless, initially similar.

[5] By this I mean not the argument to the weak conclusion that sensation terms, like any other terms, must mean the same for you as for me in order to belong to our shared language, but the argument to the stronger conclusion (which does not necessarily follow from the weak one) that when you have a given sensation you have subjective experiences similar to those I have when I have that sensation.

[6] But cf. Hacker (1971), p. 237, note 2: if he is right (which I doubt), my quotation is not to the point.

people's sensations are concerned, analogy is 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 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, fails to give a true picture of the way in which sensation language functions.

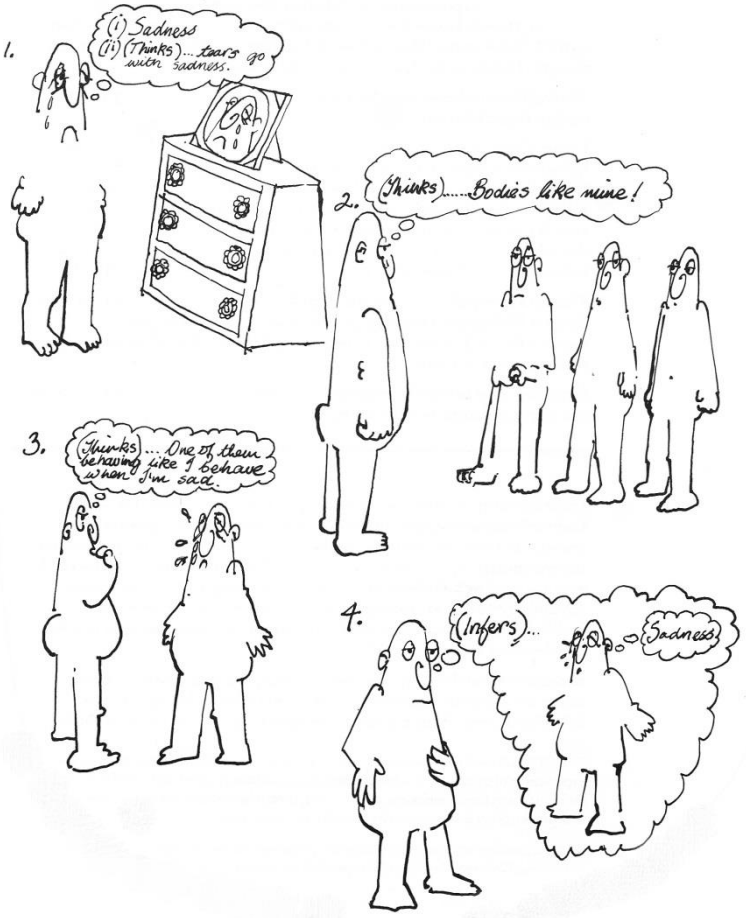

Fig. 1 *The argument from analogy*

(ii) It is argued that, if C is true, there is no way for a man to tell whether he has *x* or not, on any given occasion. Either he is under the impression that he has *x*, or he is under the impression that he doesn't have *x*. There is no way to subdivide either of these kinds of impression further into correct and incorrect impressions. For the connection between sensation and teaching links is only contingent and so one can't appeal to the **[12]** teaching links to settle the question of whether the impression is correct. So in using the term '*x*' the man would not be following a proper rule: there'd be no way of distinguishing his seeming to himself to be following a rule, and his actually following a rule. So talk of rule following, of being right and wrong, is misplaced. But one must be following rules for the use of terms if one is speaking a language. Therefore, since sensation language obviously *is* part of our language, C must be false.

Against this kind of argument there are a number of things to be said.

(*a*) Why is a check required to distinguish correct from incorrect impressions *in this sort of case* as well as in all the other sorts of case where checks are readily available? Even if the notions of correctness and incorrectness depend *in general* on the availability of checks, once the notions have been got going there is no reason why they shouldn't function quite well in particular cases, such as the case of the man under the impression that he does or does not have *x*, where *as it happens*[7] checks are not available.

(*b*) Even if checks *are* required in each individual case, why won't memory do? Perhaps *in general* memory may need public vindication: but once certified as reliable in the public arena it can take on its own private assignments without further supervision.[8]

(*c*) Even if memory were not to be trusted, there is still the argument from analogy, which can be used within the course of

[7] Even if the unavailability of checks in cases of this sort is *necessary*, a similar plea for the applicability of the notion of correctness could be made.

[8] [But we must learn the reference of '*x*' in the first place before we are ion a position to remember it.]

one person's sensational life as well as between the experiences **[13]** of two different people: intrasubjectively as well as intersubjectively. If the teaching links of *x* are present, the overwhelming likelihood is that the man's impression that he has *x* is correct, on the same sort of grounds that his assumption is likely to be correct that someone else, for whom the same teaching links are present, also has *x*.

(*d*) Even if (*a*), (*b*) and (*c*) are all rejected; even if it is said that the notion of correctness which is applicable when checks are available is inapplicable when they are not; that memory's effectiveness on the occasions when it can be checked in no way vindicates it as a device to be used on other occasions; that the argument from analogy is just inadequate; in short, that the conditions which obtain in a case where a man is under an impression about a sensation he is (under the impression) he is having are simply not good enough, and could not be good enough, for talk of knowledge, rightness, wrongness and rule following to be in order; even so there is one line of defence left. We may draw attention to facts about the use of words like 'know' and 'correct', as well as to the general way in which people use sensation language as if there were often no difficulty (as indeed there often isn't) about being sure which sensation one is having.[9] **[14]**

For example, we may say that if we find the word 'know' used naturally in cases where there is nothing but an impression to go

[9] This of course is taken by opponents of C to justify their opposition. According to C there would be a difficulty; but there is no difficulty: therefore C is wrong. But it is also possible to argue, as I do here, that the difficulty that is supposed to follow from C is illusory; that the belief that there is any difficulty depends on the imposition of a standard too high for the context of sensation. If both ways of argument are possible, why do I choose the second and apparently less straightforward one? Because the result it yields, if successful, is in accord with our intuitions about what is the case: while the result of the first way of argument is directly opposed to our intuitive views. If both ways of argument are internally coherent, that is a sufficient reason for preferring the second.

on, then that shows that an impression is sometimes an adequate support for knowledge. Perhaps when we have an impression that we are having *x* again there is no possible evidence, beyond the impression itself, that the impression is correct.[10] But if we say, nevertheless, that we know we have *x* again, that shows that no more evidence is required – not that our knowledge claim is misplaced, which is what follows from an arbitrary and rigid adherence to a more demanding notion of evidence.[11]

Besides, compare the use of terms for public objects. Here the evidence that terms are being used correctly on subsequent occasions is certainly more impressive, involving as it does such a cloud of witnesses: but logically it is deficient in a parallel way, though at one remove, as it were. For there is no independent test whereby we can verify the correctness of the collective impression of everybody that a term is being used consistently with **[15]** its earlier uses. We might all be remembering wrong, all at once.[12] But this doesn't deter us from talking of knowledge, of rule following,

[10] Actually I don't believe that we are ever this badly off, at any rate in principle: but to use the extra means which I believe is at our disposal for testing the correctness of such impressions is to presuppose both the truth of C and the truth of the identity theory of mind. To use such a device at this stage of the argument would be to beg the question in favour of the conclusion at which I am aiming: so I shall not use it, though I shall come back to the subject of the fruits of the identity theory in Chapter 5.

[11] As I have warned, the arguments in this Chapter are no more than summaries. But perhaps an extra, local warning is in order here, lest I should appear to be riding roughshod over decades of recent philosophy. Of course the theory of knowledge is a branch of philosophy with a long and increasingly technical history; of course the move of arguing from ordinary language has been subtly and lengthily discussed. Here I merely draw attention to considerations which I believe to have force: I do not pretend that the claim that they have force can be upheld as innocently of complication as this passage may seem to imply.

[12] Actually the logical side-effects of such a linguistic mass hallucination would be bound to give the game away in the end. Most terms have sufficient conceptual ramifications, interconnections with other concepts, not to be straightforwardly interchangeable with impunity. But for the sake of the argument let us allow the speculation here made.

of correctness and incorrectness – we believe that we are talking a language. Correct usage cannot be more firmly established than by discovering that everyone, after due consideration, is under the impression that the usage is correct. I do not intend by this remark to obliterate the distinction between everyone being under a correct impression and everyone being under an incorrect impression. I am not saying that certain facts about verbal usage show that we can properly describe as 'correct' both impressions which are correct and impressions which are not. No: of course the only correct impressions are the impressions which are, in fact, correct. The point was just that although there is no way, in principle or practice, to distinguish everyone's being under a correct impression from everyone's being under an incorrect impression, we can still, *pace* crude verificationism, understand the difference between these two radically different states of affairs. Why then should not a parallel point be made about sensation language? Namely that correct usage of a sensation term in a case such as the one described is established as firmly as it can be when the person who has the sensation is under the impression that he has *x* again. (A *ceteris paribus* clause, to the effect that he has learnt the use of the term '*x*' successfully in the first place, and so forth, can I hope be taken as read.) **[16]**

This need not be taken to obliterate the distinction between his being under a correct impression and his being under an incorrect impression. I am not suggesting that '*x*' is to be used to refer not only to *x* but to any sensation which a person fancies is *x*: only that it makes sense to suppose that a person is right or wrong in trusting impressions which cannot be tried before a higher court. We can understand the difference between a correct and an incorrect impression about the quality of a sensation. Indeed it is one of my main contentions that just such a distinction can and must be preserved: otherwise C is an untenable theory.

It may be argued that the cases are not parallel; that there is a difference of the following kind. In the case of the sensation we can point to a standard of correctness which is actually in use (viz. that for the use of terms about public objects) and say that the

criterion of correctness in the case of the sensation does not match up to this standard; whereas in the case of the use of terms about public objects there is no actual superior standard to which we can point. Accordingly, it might be argued, while it would indeed be silly to deny that public usage could be established as correct or incorrect (for if these notions were not allowed to apply to public usage, what *would* they apply to?), it is not silly to deny this of private usage conceived on the model of C, and hence to reject C. To which I reply, why accept only one standard of correctness, and that the best attainable? Better to face the fact that correctness about private objects (and if usage is any guide, correctness *is* an appropriate notion here) cannot match up to the same standards **[17]** as correctness about public objects. But let us grant that 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 of affairs obtains, or that whichever obtains, obtains weakly. Something is either the case, or not: if the strength of anything varies, it is that of our conviction. The fact, if it is a fact, that the tools available for testing private impressions are not as trustworthy as those available for testing collective impressions is not by itself a sufficient reason for denying that the extension of the phrase 'correct impression' includes some of the impressions tested by the inferior method. Why should the world be any respecter of our circumscribed heuristic powers? It gets right on with being thus and so, without waiting for us to find out, or troubling if we can't.

In short, the argument of this chapter has been that, even if C is true, sensation language can be both learnt and regulated. The belief that this is not possible derives from addiction to verificationist standards of language use which are not those whose satisfaction underlies our actual use of sensation language. **[18]**

2 THE IMPOSSIBILITY OF C-SUBTLE

Before I come to specifying (in Chapter 3) the version of C I favour, I shall say why I don't accept Pears's 'C-subtle' as an improvement on C-crude. In fact I regard C-subtle as in effect quite opposed to C-crude, and to that extent the presence of the element 'C' in its name makes that name a misnomer. But the two theories are linked both by being superficially similar in their specifications (closer examination will reveal the deception), and by not being behaviourism.

I begin by quoting two passages from Pears. First, a brief characterisation of C-subtle, as it applies to 'pain' (p. 152):

> the word 'pain' simply means 'sensation of a type which has such and such teaching links'. 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'.

Secondly, a defence of C-subtle against the charge that it gives sensation terms an inscrutable private reference (p. 154):

> someone who wanted to defend C-subtle would claim that, though the connection 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.

To return to *x*, then, C-subtle is the theory that '*x*' means 'sensation of a type which has such and such teaching links'.[13] **[19]** Differences between the sensations people have when the relevant teaching links are present will not be picked up in the meaning of '*x*'. If one finds C-crude unacceptable, C-subtle is supposed to provide a way of avoiding behaviourism without severing the vital links sensation language has with material object language. For meaning and reference are separated: the reference to the inner object is made in such a way that the particular phenomenal quality of this object does not contribute to the meaning of the sensation term. '*X*' is, in Kripke's phrase, a non-rigid designator: it can refer now to a sensation of one subjective quality, now to one of another quality – only the links must be right. Pears believes[14] that C-subtle is safe from the kind of attack which is fatal to C-crude.

As I have said, I do not share Pears's faith in the arguments directed against C-crude. So it is no relief to me if C-subtle is better defended against them. But even if I was in the same position as Pears, and needed a reformulation of C which avoided the weaknesses of C-crude, even then I would not turn to C-subtle. For, as I shall argue in this Chapter, C-subtle is open to the same objections as C-crude, and conversely if C-subtle is acceptable as a theory of the meaning of sensation terms, then so is C-crude, or at least the more sophisticated version of it which I shall specify in

[13] Pears uses 'teaching links' to cover 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. A link may be the natural expression of a sensation, its cause, its temporal pattern, its non-causal harbinger, or its non-expressive contemporary, to mention just a few possibilities. For fuller discussion see chapter 4.

[14] Here and elsewhere I am perhaps incautious in attributing views to philosophers. As a matter of fact, Pears does not declare his own hand, or at any rate does not declare it openly, in the chapter in question. He is concerned to disentangle issues, not to choose between theories. But of course he does occasionally point to the failure or success of this or that move in an argument. In any case, I am concerned more to characterise theories and examine them than to attribute them correctly. So for safety's sake I had better be taken to be talking of fictitious philosophers: 'Pears', 'Locke' and so forth.

Chapter 3. Both versions of C stand or fall **[20]** together. To take one step beyond behaviourism into the inner world is necessarily to go the who Le distance: there is no possibility of half measures.

In elaborating this argument I want to examine also what Don Locke (1968) has to say about C-subtle (for that is what his preferred theory amounts to), which he discusses in much greater detail than Pears does (cf. esp. pp. 96–109). Here is his summary of what he believes to be the truth about the meaning of sensation terms (pp. 100–1):

> 'pain' means 'a sensation of a certain sort', where the sort in question is determined not by how it feels, but by its causes and effects and the behaviour which characteristically accompanies it. A sensation's being a pain sensation is not a matter of how it feels, but a matter of its being of the sort caused by bodily damage and leading to pain-behaviour. Similarly a sensation's being a sensation of cold is not a matter of how it feels, but a matter of its being a sensation of the sort caused by frost and snow and leading to shivering, etc. And similarly for other sensations.

I shall concentrate on two main theoretical points which are at issue. The first point relates to the second of the two Pears quotes above. The proponents of C-subtle claim that on their theory the meaning and reference of sensation terms are separated, that is that it is not part of the meaning of a sensation term that its referent should satisfy any kind of intrinsic subjective requirement; with the result that the possible variations between people in the referents of the terms, though fatal to C-crude, now no longer affect the terms' meanings. Against this I shall argue that it is impossible to separate meaning and reference in this way; and so, if C-subtle were true, the possibility of a variable referent would have to be taken as part of the *meaning* of sensation terms, not simply as an incidental contingency.[15] This is not by **[21]** itself straightforwardly fatal to C-subtle, but it is a nail in the coffin of its plausibility.

On the second point, however, C-subtle can be more decisively attacked, so that it must be rejected whatever the plausibility of the

[15] [This needs clarifying, here and where it is recapitulated.]

conclusion I reach after discussing the first issue. C-subtle assumes that we already understand the meaning of the term 'sensation' before we come to learn the meaning of more specific sensation terms, that there is no problem about how the meaning of the term 'sensation' is taught (cf. the first passage quoted from Pears above). I argue against C-subtle that the meaning of the term 'sensation' can no more be taught (as far as the model of teaching espoused by C-subtle is concerned) than can the meanings of individual sensation terms as they are conceived by C-crude. The conclusion I reach is that, to be consistent, adherents of C-subtle must distance the term 'sensation' from subjective experience in just the same way that, in their formulations of their theory, they distance individual sensation terms from the subjectively differentiated inner referents they are commonly supposed to name: and that if they do this, the purpose for which C-subtle was adopted, name to avoid the vulnerability of C-crude without losing contact with the inner world, is barely still served.[16]

Meaning and reference[17]

Back to the first point. Can meaning and reference be separated?[18] If they can't, then it seems that the analysis of the meaning of sensation terms which must be given by a proponent of C-subtle, while not logically impossible, is yet so counterintuitive and ill-suited to the purposes of sensation language as to be unacceptable to anyone who believes himself to be expounding **[22]** the mechanics of an actual department of language.[19] This I shall shortly substantiate, by considering what Locke has to say on the subject.

Two quotations from Locke:

[16] [But C-subtle is still *possible*? (Contrast 'it must be rejected' above.)]

[17] [Omit whole section?]

[18] [Obscure locution (but cf. Pears).]

[19] [But C-subtle is already implausible in Pears's formulation, surely? He *admits* that the subjective experience can vary from case to case. In what way am I pointing to an *extra* implausibility here?]

(*a*) It is only from your own case that you know what pain is *like*; if I have never had migraine then I do not know what it is like to have migraine. But even if I have never had it, I can and do know what people are talking about when they talk about migraine, I know what 'migraine' means. (pp. 90–1)

(*b*) There are public rules which lay down when it is correct to call a sensation a pain sensation, and so it is always possible, at least in principle, for people to tell whether something is a pain sensation.[20] But this does not mean that it is always possible, even in principle, for people to tell whether such an item *exists*, whether a person does feel the sensation in question. For the argument was only that those rules which govern the use of a word and so determine its meaning must be publicly checkable, and the rule that a sensation must exist if we are to say of a person that he feels a pain is not a rule which governs the use of 'pain' in the sense of determining its meaning. To say that someone feels a pain when no sensation exists, is felt, is to say something that is false, but it is not to use 'pain' contrary to its meaning. It is not part of the meaning of 'pain' that such sensations, or any sensations at all, exist; indeed it is pretty much of a philosophical commonplace that the meaning of a term cannot determine whether the thing in question exists. 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. The argument for the need for publicly checkable rules governing the use and meaning of 'pain' does not seem to count at all against other minds scepticism. (pp. 98–9)[21]

Whatever plausibility the view expressed in the last sentence of the first quotation has, according to C-subtle (but it does not quite fit even *that* theory, according to which pain is not necessarily *like* anything), can be seen as deriving from C-subtle's claim that the particular private quality of a sensation does not determine which

[20] [Given that it exists? Even so, how to tell whether any symptoms 'go with' the sensation?]

[21] [This passage represents definite error and confusion. But are the issues separable from the ones raised by the Pears quotes? If so, the treatment of Locke could well be omitted.]

sensation it is. But if the view has any plausibility to an adherent of a version of C-crude (and as will emerge, I think it has), **[23]** this plausibility cannot be based on this partial separation of meaning and reference, which is crucially not an element in C-crude. I want to preface my examination of the claim lodged by C-subtle that meaning is untouched by reference in the case of sensation terms by showing what *is* the source of the plausibility, for C-crude, of the view referred to, so that it will be clear that such plausibility cannot. be used by C-subtle as support from the enemy camp, as it were, in its attempt (of which an example is given by the second quotation from Locke) to show that meaning and reference are separable.

Take the migraine case.[22] Suppose for the sake of argument (what I am informed is true) that the sensation people have when they have migraine are quite unlike any other sensations experienced in any other circumstances. Suppose also (what is indeed true) that I have never had a migraine. I ask what 'migraine' means. I am told that it is what can best be described as a peculiarly splitting headache, which occurs as part of a condition which can be identified by various visible and dispositional manifestations: though even to call the sensation a headache is to fail to do justice either to the acuteness of the sensation or to its lack of similarity (apart from its being painful and experienced in the head) to ordinary headaches. Do I then know what 'migraine' means? The temptation here is to say that I do. After all, I know about other sorts of pain from my own experience, and in particular about headaches; and so with a little imaginative **[24]** extrapolation I can probably manage to understand what migraine is like. Obviously, I shall understand better if I actually get around to having a migraine,

[22] It may be argued that migraine is a bad example because a great deal of medical theory is involved in the concept, making migraine less a matter of which sensation one has than a matter of which symptomatic syndrome one is diagnosed as manifesting. So I am to be understood to be using 'migraine' in what may be a somewhat deviant though popular sense (the sort of sense in which Locke seems to be using it) to designate a certain sensation or group of sensations.

but the gap between the sense I can make of 'migraine' without ever having had one, and the sense I shall be able to make of it after I *have* had one, is not large enough for us to want to say that the person who has never had a migraine does not know what 'migraine' means.

Note that I am tempted to say this for a reason which has *nothing to do* with any claim that meaning and reference are separable. I am not saying that the reason I can understand the meaning of 'migraine' without having had a migraine is that it is not part of the definition of 'migraine' that one must be having a sensation of a certain kind in order to say truly that one has a migraine. Indeed I am saying the opposite: that the subjective nature of the referent of 'migraine' is (at least partly) definitive of the sensation of migraine, but that experience of related sensations is in a case like this sufficient to give a person an adequate grasp of the nature of a referent he hasn't himself ever experienced. So to concede that we may sometimes know the meanings of sensation terms which refer to sensations we have never had in no way commits us to a theory that separates meaning and reference. It is for this reason that, although I agree with Locke in his conclusion that a person who has never had a migraine can still learn the meaning of 'migraine', I say that his view is mistaken. For he arrives at his conclusion in the wrong way, by separating meaning and reference: and so has the wrong picture of the understanding achieved by the person in question.

The misguidedness of separating meaning and reference becomes clearer if we consider a more radical case of inexperience than **[25]** never having had a migraine. Suppose now that I have never felt pain of any kind. I ask what 'pain' means. I am dished up with C-subtle's definition. Do I then understand the meaning of 'pain'? I want to say that I don't, because I don't know what pain is like. It might be argued that I do have sensations of other sorts, and I can see how people behave when they are in pain, the sorts of things that cause them pain, and so on for the other sorts of teaching links (cf. Chapter 4), so that all I have to do is to put the two together. But to this I reply that pain is sufficiently unlike

other sensations for my understanding of 'pain' before I have been in pain not to count as knowing the meaning of 'pain'. Once again, meaning and reference are not separable, and in this case lack of first-hand experience can't be made good from other sources, and so I can't properly know the meaning of the term. (Of course, this isn't an all-or-nothing affair: it can quite intelligibly be said that I partly know the meaning, 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? I could learn about the public evidence for pain as well as anyone.)

To bring out the opposition of C-crude and C-subtle in a case like this, take C-subtle's analysis of the meaning of 'pain' to be correct, and imagine that someone has learnt the meaning of 'pain' without ever having experienced the sensation to which the term refers. This is possible, for he knows the meaning of 'sensation', and that is all he requires apart from knowledge of what the relevant teaching links are. Imagine now that this person has, out of the blue, just that kind of sensation which is normally caused by a sharp metallic object penetrating the skin, only on **[26]** this occasion there is no cause of this kind. First, is this sensation a sensation of pain? But can C-subtle grant this? Suppose no pain-behaviour is elicited:[23] then C-subtle cannot grant that this is a case of pain, for there are no appropriate teaching links present, and it is its teaching links that earns a sensation its name. Even if we were to weaken C-subtle, as some would wish, so that it *could* allow that this was a case of pain, since the sensation in question is *normally* accompanied by the links definitive of pain; even so, this person cannot know that this rogue sensation is a sensation of pain (on *any* theory), for there are no teaching links by whose mediation he might be apprised of the fact. So here we would have a case of a

[23] This may seem implausible since reaction to pain appears so instinctive. But it is a moot point whether anything more than a bare minimum of pain behaviour – the so called 'natural' expression of pain – is innate: and even if it were all innate, in this case (i) the person may be so surprised that none occurs, or (ii) his innate pain-behaviour-producing mechanisms may have atrophied through long disuse.

man who (i) knows the meaning of 'pain', (ii) has a pain, and (iii) doesn't know that what he's having is a pain. This shows the emptiness of C-subtle's claim that one can know a meaning without being acquainted with the referent of the term whose meaning one knows For what good is it knowing the meaning of a term without knowing when to use it? Indeed, in such circumstances we should not normally be said to know the meaning of a word at all.[24]

So far, then, I have been discussing what C-crude would say about cases like that mentioned in the first of the pair of quotes from Locke. It should be clear that, according to C-crude, meaning and **[27]** reference cannot, even in support of a view like this one of Locke's, be prised apart even to the extent of saying that, though a sensation must be present for the term to be applicable, the private nature of this sensation is immaterial. C-crude lends no support to C-subtle on this score.

On, then, to C-subtle, which already looks an implausible enough theory *even before* its claim that it succeeds in separating meaning and reference is considered: even the move of making the particular subjective quality of the referent irrelevant puts it beyond the pale from the point of view of C-crude. The point of the extra claim is to show that C-subtle modifies C in such a way that it can be cleared of the charge that, because it is possible that the same name refers to subjectively different sensations in different people, and even perhaps to subjectively different sensations on different occasions in the same person, so the meaning of sensation terms may be different for different people, may even change for the same person. This charge could be completely dismissed if it could be shown that the meaning of sensation terms is not only no

[24] I do not wish to maintain the general thesis that a person knows the meaning of a word only if he knows both the truth conditions of statements in which it features and how to recognise whether those conditions obtain. Such a thesis would be false (cf. e.g. Putnam (1973)). But in some cases something very like it seems to me to be true: and the case of 'pain' is one such.

respecter of subjective variation,[25] but is simply nothing to do with the terms' reference. And so C-subtle aims to show this. But is this aim achieved? If C-subtle is true, Pears says (p. 154), the private reference '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".' But *is* it so made? Agreed, it is no part of the meaning of the sensation term, on this model, that just *this* (phenomenally differentiated) type of sensation is involved: but it *does* seem to be part of the meaning that *some sensation or other* is involved; just as the referent of the phrase 'a person's **[28]** favourite cocktail' (which Pears uses as an illustration: cf. p. **15**), must, after all, be at least *a cocktail* (and not, for example, a plate of egg and chips). But even *this* degree of connection between reference and meaning seems not to be conceded. What is more, Locke seems at times concerned explicitly to deny it. (Pears may just be talking loosely, and only mean the lesser degree of separation.) Locke seems astonishingly to want to say that, for a given sensation term to be applicable, it isn't simply that it doesn't matter *what* private event is going on – it doesn't even matter whether *any* such event is occurring. This I reject. The second of the two quotations from Locke above (p. **19**) shows the counterintuitive consequences of espousing a separation as excessive as this. It is a strange passage in more than one respect:

(*a*) Locke's remark (pp. 98–9) that 'it is not part of the meaning of "pain" that such sensations, or any sensations at all, exist' contradicts, at least on one interpretation, his view expressed (e.g.) on p. 100 that' 'pain means "a sensation of a certain sort" '. Cf. also his assertion on p. 109: 'It does matter that some object exist.' Of course, if we take the remark to amount merely to the claim that it does not follow from the fact that 'pain' means 'a sensation of a certain sort' that any sensations of that sort exist, any more than it follows from the meaning of 'cow' that any creatures of that sort

[25] [But this first step is sufficient, surely? The complete detachment of reference is rather a red herring.]

exist, then perhaps[26] **[29]** what he says is true.[27] Indeed this way of taking the remark is suggested by the rest of the sentence from which it is taken ('indeed it is pretty much of a philosophical commonplace that the meaning of a term cannot determine whether the thing in question exists'), for this could not serve as an amplification of the (false) point that the presence of a sensation is immaterial when there is a question of whether somebody is in pain or not, which is the other way of interpreting the remark in question. But if we do take this remark of Locke in the way in which this sentence suggests, as a point about the impossibility of deriving existence claims from meanings, then it is nothing to the purpose, which is to show that meaning and reference are

[26] Though it would be at least very odd to pick out a group of publicly observable phenomena as definitive of a sensation if no sensation at all, or no regularly similar sensation, had ever been known to occur in company with those phenomena. This is certainly a contingent matter, but sufficient, for all that, to stop us in practice from providing an appropriate word. Would a term defined as 'the sensation one has at 5 p.m. on one's thirtieth birthday' have a sense? I suppose it would, but no one would ever ordinarily invent such a term.

[27] Though I can't see that the stronger claim is true, viz. that it doesn't follow that any sensations at all exist: for how are we to give sense to 'sensation' if no sensations exist – or for that matter to 'animal' if no animals exist (if that is thought possible, read 'if no physical objects exist')? 'That which there is in such and such situations over and above publicly observable phenomena' won't do as a definition of 'sensation', for *ex hypothesi* nothing exists to satisfy this description, so that the description fails to specify a class: it amounts to 'that which there is over and above that which there is', which is nonsense.

It is no counterargument that we can give a sense to 'unicorn', even though (by definition, if Kripke is right that unicorns are essentially fictional) no unicorns exist in the real world: for the sense we give to 'unicorn' is parasitic on the sense we give to a range of terms which name real animals such as 'cow'. Given an understanding of terms which refer to real entities, we can by analogy understand terms which refer to fictional entities. But the imagined case of 'sensation' in a sensationless world is not of this kind, for the proffered definition of 'sensation' makes sensations not fictional entities analogous to real ones, but real non-entities, which is to say, inconceivable.

Of course, given that there *are* sensations, one might be able to specify them negatively, by exclusion, as it were, in terms of public phenomena (see further pp. **32**f.): but this, it should be clear, is not what is being suggested here.

separable.[28] **[30]** To say that the meanings of terms don't entail the existence of entities is not the same as to say that it is not one of the necessary conditions for the successful use of a referential term to refer – that is, part of that term's meaning – that something of a specified kind should exist for the term to refer to. A referential term (and even according to C-subtle sensation terms are referential) can certainly have a meaning without there being any appropriate referents for it: but if it is successfully used in a referring role, then necessarily an appropriate referent exists.

(*b*) '[T]he rule that a sensation must exist if we are to say of someone that he feels a pain is not a rule which governs the use of 'pain' in the sense of determining its meaning' (p. 98). On what grounds is this rule segregated from the other rules governing the use of 'pain', rules which *do* determine its meaning? These other rules will presumably mention a set of conditions for the ascription of pain, and specify that the satisfaction of a suitable subset of these conditions, or of a prominent enough individual condition, will be sufficient to establish that pain is to be ascribed. Examples of such conditions might be 'He has just had a gallon of molten lead poured over his head / He is hopping round on one leg, screaming and grasping his other foot / You have just punched him violently in the stomach' etc. ad infinitum (whether or not it is reasonable, the specification of pain in this way via teaching links is a hopeless task). But since the meaning of 'pain' is said to be '*sensation* with such and such teaching links', the presence of a sensation is a necessary condition for the correct ascription of pain, and by parity of reasoning the requirement that one be present is one of the rules determining the meaning of pain. **[31]**

What *can* Locke be after here? It sounds so crazy. If I said 'It's true that he's in pain but I don't expect he feels a thing', wouldn't that be *just* the sort of remark that showed I didn't know the meaning of 'pain'? The way Locke expresses himself shows that he is *not* simply making the point that meanings are not grounds for existence claims. Nor, surely, can he be making the point that one

[28] [Yes – and perhaps this is all that needed saying.]

can use a term *correctly*, in accordance with the rules governing its use, without succeeding in making a true statement: this can be done in various ways, but in particular one can assert that someone is in pain, believing that he has a sensation of the appropriate kind, when in fact he hasn't. So one might make the point that the test of whether someone is obeying the rules governing a term's meaning is to see whether, when he uses the term, his beliefs are of an appropriate kind: and beliefs, notoriously, don't necessarily match facts. So one can't demand that, if we are to say of someone that he feels a pain, a sensation must, simply, *exist*: rather the rule is that a sensation must be *believed to exist* (by us). But in the first place it is easy to make it clear that this point is being made: it is not a difficult point to express, and there is no sign that it is the point Locke is concerned to make. And, secondly, it is not a point which, when made, successfully achieves the aim of separating meaning and reference. For it is always possible to replace a statement of a necessary condition for a term's use, say 'If *x*, then necessarily *y*', by a statement of a necessary condition for the coherence of a belief whose statement involves the use of that term, thus: 'If *x* is believed to be the case, then necessarily *y* is believed to be the case.' But such a restatement in terms of beliefs does not show that it is not part of the meaning of *x* that *y* **[32]** (whether *y* is ever *in fact* the case or not). We can say that we only say truly that someone is in pain if he is having a sensation of an appropriate kind: or we can say that we only understand the meaning of 'pain' if, when we assert that someone is in pain, we believe that he is having a sensation of an appropriate kind. But on either account the rule that a sensation must exist if we are to say of someone that he feels a pain is a rule which governs the use of 'pain' in the sense of determining its meaning. There is apparently no alternative to taking Locke to mean what he says. In which case his point is straightforwardly false.

(*c*) 'To say that someone feels a pain when no sensation exists, is felt, is to say something that is false, but it is not to use "pain" contrary to its meaning.'

Not necessarily. Locke is oversimplifying. There are two kinds of mistake sensation ascribers can make: they can fail to know the facts, and so assert a falsehood, even though they know the meanings of the words in which they do so; or they can fail to know a word's meaning, and so say something they don't intend to say, even though they know the facts. And there are two ways in which someone can fail to have a specified sensation: either he has no sensation at all, or he does have a sensation, but, because of the links with which it occurs (I speak in C-subtle terms), it is not the sensation specified.

Locke Simply pairs asserting a falsehood with the absence of sensation, and ignorance of meaning, misuse of language, with the presence of the wrong teaching links. But one can assert a falsehood also because one has mistaken the links; and the reason one says that someone has *x* when he in fact has no sensation at all can be also that one does not know that it is part of the meaning of *x* **[33]** that *x* is a sensation. In brief, whether the mistake arises because there is no sensation, or because the wrong teaching links are present, two kinds of mistake are possible: not knowing a word's meaning, and not being correctly apprised of the (non-linguistic) facts. For practical purposes these two kinds of mistake may amount to the same thing: ascriptions made under their influence will be false. But the reasons for the two kinds of mistake can certainly be distinguished.

Now consider the distinction Locke attempts to draw in this sentence (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.'[29] Well, the distinction between the two kinds of rule we may

[29] Surely Locke is not here relying on the trivial point that the quoted rule, to be a meaning rule, should strictly be expressed in the formal mode, for example thus: 'The word "pain" is properly used to ascribe a property to a person only if that person has a sensation', or even more formally: 'The word "pain" means "sensation of such and such a kind" '? For meaning rules are often, and unmisleadingly, stated in the material mode. Similarly, when in the

grant, for all that they amount to the same. But, given the distinction, Locke's classification of the rule he quotes as an example exclusively of the second kind of rule is a mistake. For the quoted rule can be a rule of either kind, and can be broken when it is wearing either hat. Which hat it is wearing in a particular case of its contravention is determined by whether it is the speaker's beliefs or the way things are which fail to fit the assertion in question.

The kind of case in which the broken rule *is* a rule of the second kind, a rule determining when it is true to say of someone **[34]** that he feels a pain, is precisely not the kind of case to use to show the separability of meaning and reference. All it shows is that it is possible to be mistaken about the facts. Once again, then, the attempt to show that the reference of a sensation term plays no role in its meaning is a failure.

I take it, then, that meaning and reference are not separable even by the adoption of C-subtle, let alone according to C-crude. Naturally, since I espouse C-crude, I believe both (i) that acquaintance, direct or analogical, with the referent of a sensation term is necessary in order fully to understand the term's meaning; and (ii) that one of the rules for the use of the term to describe the world correctly, and so part of the term's meaning, is that a suitable private referent be present. This is not the same as saying that any such referents need actually exist for there to be a term with such a meaning: this may or may not be true (because of (i), I believe it must be true, at any rate if a term with such a meaning is to be understood), but it is a different claim.

Of course, it is true that C-subtle analyses the meaning of sensation terms in such a way that the role of the referent in the meaning of such a term is pruned down to the absolute minimum. By defining '*x*' as 'sensation of a type which has such and such

passage quoted from Locke in the previous paragraph he talks of saying 'that someone feels a pain when no sensation exists', it does not follow from the fact that there is no explicit mention of words that it cannot be the meaning of a word that is under discussion.

teaching links', one's requirement of the referent is reduced from the requirement that there be a sensation of a particular subjectively differentiated kind (for a sensation is now to be classified by its links, not by its inner quality) to the requirement that, simply, *some* sensation be present, no matter what its subjective nature. Such a move is possible, but:

(*a*) The rules which determine the meaning of a sensation term still require, albeit it is a minimal requirement, that there be *a* **[35]** sensation for the term to refer to. Reference and meaning have not been entirely separated.[30]

(*b*) Such a model of sensation language fails to do justice to one of the main *points* of the sensation language game. Our belief that we can communicate about the qualities of our sensations is illusory if C-subtle is true. It will in effect be part of the meaning of sensation terms that the private nature of the referents of a given term need not be constant (though for C-subtle the ascription of private natures is unintelligible).

(*c*) In any case, because C-subtle makes use of the notion of a sensation, it is vulnerable to the argument presented in the following section. A minimal reference to the inner world is as bad as the greater involvement of C-crude. So the claim made out by the defender of C-subtle in the second quote from Pears on p. **15** doesn't stick. But this is to anticipate.

The meaning of 'sensation'

Given, then, that even C-subtle is committed to saying that sensation terms refer to sensations (even though it doesn't matter what the sensations are subjectively like), and that it is part of their meaning that they do this, we can come to the second main point at issue between the proponents of C-subtle and myself.

C-subtle assumes that the notion 'sensation' is understood. But how, if there must be public criteria for the correct use of a term (for this is an essential part of the view of the proponents of C-

[30] [Pears doesn't want total separation, I guess.]

subtle), is the meaning of the term 'sensation' to be taught and its use monitored? There is no way of telling, on any particular occasion or in general, whether another person is having a sensation (no matter of what kind) any more than there is a way of telling whether someone else is having qualitatively the same kind of sensation when he says he has sensation *x* as I am having when I say I am having sensation *x*. And it follows from this that there are no **[36]** paradigmatic situations in which it is known that sensations occur, which can serve as the foundations upon which to build a language learner's understanding of the meaning of 'sensation'. It's no good saying we learn 'sensation' as the general term for all the specific sensations, rather as we perhaps learn 'material object' as the general term for all the 'moderate-sized specimens of dry goods'[31] that furnish the world: for we need to understand the meaning of a whole range of specific terms before we can learn the meaning of a general term on their basis, and according to C-subtle we can't learn the meaning of specific sensation terms except as 'sensation of a type which …'. But *this* presupposes that we already understand the meaning of the general term whose meaning, it is proposed, we shall come to understand on the basis of an understanding of the meaning of a number of the specific terms subsumed under it. The learning situation is quite circular.

Since there are no paradigmatic situations in which sensations can be said to occur, it follows that 'sensation' cannot be a genuine word in the language, for there are no publicly checkable rules for its correct use. So C-subtle is open to just the same kind of objection as C-crude, viz. that its model of the meaning of the term 'sensation' produces the absurd result that the meaning could not be learnt; and so the model must be wrong. If '*x*' can't be the name of a subjectively differentiated type of sensation, nor can 'sensation' be the name of a subjectively differentiated type of experience.

[31] Austin (1962), p. 8.

Naturally, I do not hold with this argument, since I believe in the argument from analogy. The teaching of 'sensation' by analogy is even easier than the teaching of some specific sensation term. We can surely guarantee to produce *some* sort of sensation **[37]** in a language pupil as a basis for teaching him the use of 'sensation'. But it is precisely the assumption which lies behind the belief that this is the way in which sensation language is learnt that is rejected by the opponents of C – the assumption that publicly observable events are in any way reliable as signs of particular kinds of inner event, or of any inner event at all. So for the time being I am proceeding on the assumption that the argument from analogy is unacceptable, in order to show that without it there is no satisfactory way of implicating inner events in the language game played about sensations. C-subtle must not be allowed to deny the use of this argument to C-crude, only to make surreptitious use of it on its own behalf.

So the teaching of the word 'sensation' is beset by precisely the same difficulties as the teaching of individual sensation words. If the proponents of C-subtle are to be consistently opposed to the supposed weaknesses of C-crude they must give an analysis of the meaning of 'sensation' analogous to the analysis they give of the meaning of particular sensation terms. Just as '*x*' is analysed as 'sensation with such and such teaching links', so 'sensation' will have to be given an analysis, as a rough first shot, such as 'thing with such and such teaching links', where the teaching links which are specified in this analysis will comprise the set of all the specific teaching links mentioned in the analysis of particular sensation term such as '*x*'.

Once such an analogous analysis has been given, then just the same sorts of points which are made by the supporters of C-subtle about the irrelevance to the meaning of specific sensation terms of **[38]** possible variation from person to person of the subjective experience which accompanies the teaching links definitive of those sensations can also be made about the general term 'sensation' itself. To adapt my first quotation from Pears on p. **15**:

the word 'sensation' simply means 'thing of a type which has such and such teaching links'. So if, unknown to us, and perhaps unknowably to us, the same teaching links pointed to different kinds of item in different people, those differences would not be picked up in the meaning of the term 'sensation'.

It might be argued that the word 'thing' in this adapted quotation is too general. I could have specified the type of thing in question more narrowly by writing 'experience'. After all, not everything counts as an experience, while nothing is disqualified from being a thing – so to write 'experience' is to limit the field within which the required item is to be found, whereas to write 'thing' is to give no positive clue at all as to where to look: as long as the right teaching links are present, anything which is present in addition would count as the sensation. But this isn't right: a sensation must surely be at least some sort of *experience*?

This objection seems plausible enough until it is seen that it only delays for one more round the eventual inevitable confrontation with a purely negative characterisation. If in the original Pears quote one substituted for 'pain' not 'sensation' but 'experience' (experiences being the genus of which sensations are a species), then what. would take the place of 'sensation' in the original analysis? Event? Thing? Phenomenon? Object? State? Eventually it is necessary to acknowledge that no positive subjective characteristic can be specified if the arguments against C-crude **[39]** are successful.

Perhaps the proponents of C-subtle would be willing to accept these consequences. But if they do accept them, it seems that the connection with the inner world, which it was one of the aims of C-subtle to protect, is getting perilously flimsy, in the following way. When we were forced to abandon C-crude for C-subtle, although we had reluctantly to give up the idea that my sensation x was subjectively like your sensation x – or at any rate we had to accept that whether the two were subjectively similar had nothing to do with the meaning of 'x' – at least we could comfort ourselves with the thought that both of us had a *sensation* at least. That

seemed a safe enough assumption. We know what sensations are like, and at least the other person had one of *those*. But now even this solace is denied us. 'Sensation' is not to be a class name for a set of experiences whose nature we are acquainted with only from our own case (notice the constant parallels that crop up in the case of individual sensation terms), for as such its meaning could not be learnt. Rather it is to mean (waiving the delaying tactic with 'experience') 'whatever there is over and above any set of teaching links which are definitive (presumably, only when there *is* something over and above them) of the presence of a sensation (type unspecified)'. 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'.

What is now left of our foothold in the inner world? All we are entitled to claim, it seems, when we use the term 'sensation' (or, analogously, any other term naming a species of the genus subjective experience), is that something non-outer, of an unspecified and unspecifiable subjective nature, is in the offing, and is to **[40]** be defined purely by reference to its observable accompaniments. Is this a claim worth preserving? I think not. For if the inner is characterised solely by not being outer, and a particular type of inner event by not being a particular type of outer event (though its contemporary), then the inner has no positive characteristics left at all, as I have noticed above. A purely negatively characterised world of subjective experience might as well be abandoned.[32] Mere existence is an unappealing attribute, especially for something whose identity changes in step with the company it keeps.

The extreme bloodlessness of sensation terms as viewed by C-subtle can be brought out by looking at what Locke says about the

[32] Of course abandonment is not required logically. But I am suggesting that the only reason adherents of C-subtle continue to adhere is that they fail to see that their theory has consequences quite as far-reaching as those discussed here. If they did see this, they might think again.

case of the beetle in the box.[33] He claims that learning the meaning of 'beetle' in the imaginary case is parallel to learning the meaning of 'pain' (or, presumably, of 'sensation') in the actual case, in various ways. I shall dispute this.

Locke says that Wittgenstein is wrong to say 'the box might even be empty': the object drops out of consideration as irrelevant only in that it does not matter what, the object is *like in itself*, not in that it does not matter whether it *exists*, for this *does* matter. And people know what it is for one thing to be inside another on the basis of their experience with cases exposed to public view, and 'so if a person has nothing in his box he will tell us' (p. 108). 'Similarly if a person did not feel pain, then when he **[41]** is told that by "pain" (my quotes) we mean not his behaviour, nor what he says, nor the part of his body where we have kicked him, he will naturally retort "But there isn't anything over and above these – I am not aware of any pain." '

Just as, then, 'beetle', is to mean 'that thing, whatever it is, which is in the box', so 'sensation' is to mean 'that thing, whatever it is, which is present in situations of such and such kinds over and above publicly observable phenomena'. On the surface, this parallelism is sound: but the grammatical similarities between the two analyses, particularly the two occurrences of 'that thing which', amount to a subterfuge which glosses over a disanalogy between beetles and sensations, a disanalogy which, if acknowledged, makes the analysis of 'sensation' look intolerably thin.

A beetle in a box is an ordinary sort of thing which just happens to be screened from public inspection in this imaginary example. We already know what physical objects are on the basis of unimpeachably public instruction, so when the question arises of whether or not somebody has a beetle in his box, everybody knows what sort of an item it is which is being asked about. But the case of a sensation is different: a sensation is (necessarily) *never* publicly observable; it is hardly a sort of material object; indeed it is called an object at all only by a considerable stretching of the term's

[33] Wittgenstein (1953), § 293.

meaning. A sensation is pretty much *sui generis*. So when the question arises of whether or not somebody is having a sensation – a question which must arise, and be answered, if we are to know whether a sensation term is being correctly used – it is not the case that everybody knows what sort of an item it is which is being asked about. We cannot appeal to an already **[42]** understood notion of material object, or to an already understood notion of one thing being *in* another, in order to give some character to the item whose presence is in question. All we can say is that there is (or isn't) something there in given company at a given time which is not publicly observable. Even the use of the notion 'thing' is less informative than it might seem. Sensations are not remotely like any other 'things'. So to call them 'things' is to characterise them hardly all, beyond giving them the minimal properties of spatiotemporal thinghood, such as duration and location. The fact is that any acquaintance which we have with the positive characteristics of sensations must come from our own personal experience of sensations, and cannot be provided by purely external tuition. So if 'sensation' is to be a positive notion at all, we will have to rely on the argument from analogy to guarantee that its positive elements are intersubjectively comparable and intrasubjectively constant. And, as C-subtle correctly points out, if we can't provide such a guarantee, there is no point in claiming the positive characteristics: they would be wheels that turned without being part of the mechanism.[34] Their mechanical role can be assured only be requiring that sensations have a certain subjective nature, on pain of losing their identity.

It would, as I have said, be a stalling manoeuvre to argue that the term 'sensation' could be introduced via already understood locutions like 'subjective experience', 'what you feel', 'whatever you are directly aware of' and so forth. The learning of the meaning of these locutions is beset by precisely parallel difficulties. Without

[34] [Cf. Wittgenstein (1953), §271: 'a wheel that can be turned though nothing else moves with it, is not part of the mechanism'.]

the argument from analogy, we can get no positive foothold of any kind in the inner world.

So those who reject the argument from analogy should not allow themselves to derive any consolation from C-subtle. If particular **[43]** sensation terms cannot be tied to particular subjective experiences, then neither can 'sensation' be tied to a particular sub-genus of subjective experience: nor can 'experience' be tied to a particular kind of anything – it can only be excluded from other kinds of things. To be fair, a purely negative characterisation is not the same as no characterisation at all: in a limited context it can serve quite well (in the context of men and women, one might quite satisfactorily characterise women as being not-men). But in the wider context imagined here it comes very near it: near enough to make the inner world a totally uninteresting parasite on the outer one, describable only in terms of its borrowed clothing, with its inherent nature barred from discussion as being unspeakable (it cannot even be supposed to exist unspoken). In this predicament it is better to cut one's losses and accept that for any non-behaviourist who wishes to prescribe anything other than an entirely etiolated additive to behaviour as constitutive of mental states, the argument from analogy, or some device with an equivalent effect, is a necessity.

The attack on the acceptability of C-subtle which I am mounting is as much a psychological one as a philosophical one. C-subtle is a theory which continues to radiate an attractiveness based on certain features it has, long after philosophical arguments against those features have been heard and accepted. So I have some more to say about C-subtle. If it is thought that I am overdoing it, I can quote Ryle: 'Some readers may think that my tone of voice is excessively polemical. It may comfort them to know that the assumptions against which I exhibit most heat are assumptions of which I myself have been a victim. […] I am trying to get some disorders out of my own system.'[35] **[44]**

[35] Ryle (1949).

But the rest of what I have to say about C-subtle must come after I have spelt out my preferred version of C-crude, for it depends on a distinction which is essential to that preferred version. Which is to say, it must wait until later in the next chapter. **[45]**

3 C-CRUDE-SUBTLE

If, then, C-subtle is right about the vulnerability of C-crude, we have now shown that C-subtle is vulnerable in the same way. But once again, given that I do not accept the arguments against C-crude, nor do I accept the force of similar arguments directed against C-subtle. Either theory is logically possible: it is open to us *either* to require that a given sensation term refers to a given type of subjective experience, *or* merely to require that it refer to some member of a given species of subjective experiences, no matter which member. Which option do I prefer? As I have said, our intuitions point to the first. I prefer C-crude.

But C-crude needs further modification before I will unreservedly champion it. And so, at last, I come to the specification of my own version of C, which I do by modifying in two important ways the crude version of the theory quoted from Pears on p. **5** above. Having made these modifications, I call my theory 'C-crude-subtle': this name is meant to be understood as 'the subtle version of C-crude' (let us call the Pears version 'C-crude-crude'), *not* as 'a mixture of C-crude and C-subtle', which is precisely what it is not supposed to be.

The first modification is to concede the importance of teaching links for learning the meaning of sensation terms, but to introduce them into the theory in quite a different role from that in which they were introduced into C-subtle. In order to state this role clearly, I use certain semi-technical terminology which can be found defined in Kripke (1972). I hold that '*x*' means 'sensation of *this* type', where, just as for C-crude-crude, the type is subjectively differentiated; and that the reference of 'this' is fixed by **[46]** teaching links which characteristically, though contingently,

accompany the sensation in question. It is not part of the *meaning* of '*x*' that, if a person has *x*, these teaching links are present: they are simply contingent features of the state of people who are having sensation *x*, which is a particular phenomenally differentiated kind of subjective experience rigidly designated by the term '*x*'. Of course the analysis 'sensation of *this* type' is only of any practical pedagogic use in a learning situation where ostension of the sensation via its teaching links is possible. If one wants to cater for less direct learning situations one will have to analyse '*x*' more circuitously as 'sensation of a particular subjective quality whose identity is taught via, but not defined as such by, the fact that it characteristically has such and such teaching links'. I preferred to give the ostension-analysis first because it is starker, and less easily misunderstood as version of C-subtle. The second analysis can seem like a version of C-subtle because the teaching links are mentioned, apparently, in the definition. But this is an illusion: the teaching links are *in no way* definitive of the sensation. The only reason they occur in the second analysis is that, since the only way it is possible to pick out particular sensations as the ones being talked about is to mention their accidental public symptoms, a definition which is to be of any practical use must have a rider telling the learner where to look for the item being defined. 'Sensation of a particular subjective quality', though a perfectly correct definition, is incomplete in that it does not distinguish one sensation from any other. Unfortunately, in the case of sensations, it is not possible to complete the definition by listing any uniquely **[47]** identifying properties of the definiendum. So one has to give, as it were, a foster definition in terms of contingent symptoms. But just as a foster child's existence is in no way dependent on the existence at any time of its foster parents, so there is no necessity that a sensation ever had the teaching links it in fact has. Without them, or some others, the use of a term referring to it could not have been taught: but the sensation could still have existed, and

could have felt just the same. A foster definition defines a term in terms of non-necessary characteristics of its referent.[36]

C-crude-subtle, then, is not to be taken as a version of C-subtle. The crucial difference is that *x* is not to be identified as *x* because certain teaching links are present which are logically sufficient conditions for a sensation's being *x*; but it is to be identified as *x* because it is *x*. The presence of the links may be strong evidence that we have a case of *x* (if it is another person's sensation, perhaps the only evidence), but that is all. Sometimes *x* may occur in the absence of the teaching links[37] which are present in those paradigmatic situations where the use of sensation terms is learnt. Still we can say we have a case of *x* (C-subtle cannot).

[48] Given that the subjective nature of the referent is definitive of the sensation, it becomes impossible to subscribe to a remark like this one of Locke's (p. 98): 'If a person sincerely asserts that he feels a pain when in fact he is laughing, smiling and showing no signs of strain or injury, then we know that he is misusing the word "pain".'

Quite apart from the possibility that he may be acting, his sensations may have become mismatched with his behaviour: but as long as they were normally matched when he *learnt* the word, we may believe him. Admittedly this sort of case is far from straightforward, and our belief that this person is in pain, even if supported by evidence beyond his own avowal – e.g. neurophysiological evidence – would not be an unclouded belief: for the normal teaching links of a sensation *do* acquire a strong

[36] Perhaps very many definitions are of this kind, if it is true that necessary characteristics are seldom those obvious ones in terms of which everyday 'definitions' are usually framed. But presumably such 'definitions' are to be shunned in favour of natural definitions where possible, in the interests of rigour, just as one does not foster a child unless its natural parents are inadequate.

[37] More likely perhaps than in the case of other foster-defined words. If Kripke is right, yellowness is not a necessary property of gold. But it would be surely be more upsetting to find that gold (or even some gold) was in fact blue, than to find cases of, say, pain in the absence of characteristic teaching links.

psychological association with the quality of the sensation they usually accompany, or even a logical connection, so that, if the partnership were to change, the meaning of the relevant sensation term would to some degree be affected. This kind of consideration may well be urged against what I have been saying about C-crude-subtle. I admit that the picture of the meaning of sensation terms which I have been sponsoring is idealised, in that there may be no actual sensation term whose meaning is determined *purely* subjectively: we are practical people, and it would be surprising if a term of this kind could survive for long. But I thought it worth committing the idealisation for the sake of clarity: it is sufficient if terms of this unsullied kind are at least *conceivable*, for the claim of the opponents of C-crude is not simply that they don't exist, but that they couldn't. I expand this point later. Here let me just say that, even if teaching links have got entangled in the meanings of all sensation terms, even this degree **[49]** of involvement is a far cry from the claim (essential to Pears's and Locke's versions of C-subtle) that teaching links constitute the only criteria for the use of a term. Subjective quality may still be the *dominant* criterion, if not the only one. So that if nature nodded and teaching links and subjective quality parted company, the term would stick to the subjective quality and desert the links, albeit this would mean a slight change of meaning. It may be true that teaching links of a certain kind are among the essential criteria for a situation being appropriate for *teaching* the use of a term (later, perhaps neurophysiological teaching methods may make them redundant): but learning and use must be carefully distinguished. Once a term has been learnt, it can be attached by the successful pupil to the relevant subjective experience, and detached from its paradigmatic teaching links – either through desuetude, or by force of circumstances (a rare or non-existent contingency). So when Locke says, in the next sentence after the passage quoted on p. **41**, 'There are public rules which lay down when it is correct to call a sensation a pain sensation', he might more truly have written: 'There are

public guidelines[38] which enable one to learn correctly which sensation to call pain sensation.'

Considerations of this kind, to do with the involvement of teaching links in meanings, are not unconnected with the second main modification I wish to make to C-crude-crude, to which I now come. According to C-crude-crude, there is an instance by instance **[50]** necessity that, if a sensation term is to be properly used, then just the right kind of sensation must be present on that particular occasion: just as, whenever the word 'rose' is used properly of an object, that object must be a rose. I am prepared to soften this requirement a little for sensation language – a dangerous ploy, likely perhaps to be the thin end of some criteriological wedge: but I think one *can* get a foothold part way down the slippery slope that leads in that direction.

First I notice the inevitable fact that since we don't introspect the feelings of others, and since our aim in using sensation language is often to decide on a course of action (sympathy, medication, punishment, congratulation, advice etc.) rather than to compare phenomenal notes, the teaching links which make it possible to learn sensation language will usually be prominent in our minds as we perform the action, and may indeed seem to a dispassionate behaviouristically-minded observer to be quite sufficient to explain all human linguistic behaviour about sensations. Sensation terms may seem like shorthand ways of referring to typical groups of public phenomena. This is only natural, given that sensations, being private, are dealt with by this second-hand method of outward signs or links.

Having noticed this fact, I concede that there *will* be cases where the discovery that the wrong sort of sensation is in the offing would not bring us up too short: the action we choose would be the same. Perhaps the links play a sufficiently autonomous role in the meaning of the relevant term for variation in the sensation to

[38] These guidelines might change over time if normal matchings between sensation quality and links changed: but the meaning of sensation terms would not be much affected thereby if the dominant criterion remained subjective.

be relatively unimportant. The dentist who is using pain as an indication of which teeth need repairing is not going to be worried if he discovers that there is a huge phenomenal variation in the **[51]** sensations his patients have when they complain of toothache.[39] This is an extreme case, coming at the links end of a continuum of cases where the relative importance of the links and the quality of the sensation varies: not all sensation terms are alike, nor are all uses of a given sensation term – the purposes of remarks in sensation language are various.

But that is all that I will concede. Although in a particular case the specific quality of the sensation may not matter, it is a vital background assumption of our sensation language that correlations, intra- and inter-subjectively, occur between sensations and teaching links. The assumption may be false, but it is none the less built in to our linguistic habits, and would not be cheerfully shed.

In my specification of C-crude-subtle I have made it a logical requirement, part of the meaning of sensation terms, that a sensation of a given kind should be of a definite subjective quality. But it must be conceded. to the opponents of forms of C-crude that this kind of logical requirement differs in an important way from most logical requirements, just because it is private. That is why I refer to it cautiously as an assumption. I have said that this assumption is vital: but there is a sense in which to characterise the assumption as a straightforward logical requirement is misleading. If the requirement were not fulfilled **[52]** (in certain ways), the mistakenness of the assumption would not show up. But if the assumption were straightforwardly logically vital, its falsity would entail the incoherence of sensation language. So it does, strictly speaking: for if the assumption is false, then C-crude-subtle's view of the meaning of sensation terms must be wrong – terms with meaning rules of that kind would have no application. The trouble

[39] It may be argued that these sensations must be at least unpleasant, otherwise why bother to repair the teeth? But rotten teeth can spread deleterious infection as well as giving pain. The dentist will be more interested in the sensation's cause than in its quality.

is that C-subtle is always ready to step into the breach. And although under its aegis the analysis of the meanings of sensation terms would have to differ, the actual mechanics of its learning and use might be untouched[40] (particularly if we imagine that, though the vital assumption is false, the parties to it still regard it as true). As far as outward and visible signs went, there would be no change.

So the private kind of rule which is such an important feature of C-crude-subtle is unique in that, if it is broken, there need be no public repercussions. And this is precisely what leads opponents of C-crude to insist that it is not really a rule at all. Now I do not want the qualifications for rulehood to be as constricting as this; otherwise I should not have incorporated a private rule into my preferred theory of sensation language. But in deference to the view that a rule which has no public effect is no rule, I shall now try to show that a requirement which does not achieve the full status of a rule can yet be as significant **[53]** for the health of a concept as ordinary meaning rules. So that even if I were forced to call such a requirement 'psychological', say, rather than 'logical', a case could still be made out for the inadequacy of C-subtle: the theory would then be inadequate not because it got the rules wrong, but because it was framed as if the rules were all that mattered. This may look as if it boils down to an issue of mere nomenclature: but there is such venomous philosophical opposition to the idea of private rules that I think it is worth showing how, even if such opposition is yielded to, it is still necessary to admit the force, albeit under another guise, of the considerations which led to the framing of the private rules in the first place.

So now we are to say that the assumption about intersubjectively similar subjective quality which I am championing is one of

[40] That this is possible shows, according to C-subtle, that there was no content to C-crude's claim that there is something over and above these mechanics that matters for meaning: if you still do just the same things, what does the claim that something has changed amount to? The answer which I am trying to defend is that something that does not glitter may yet be gold.

those conditions for the functioning of a concept or set of concepts whose non-fulfilment would be practically disastrous, even though its role is not immediately visible in the surface logic on any particular occasion of the use of one of the concepts. Indeed the surface logic of sensation concepts is compatible with the rival theory (C-subtle) whose truth I am claiming would be fatal to sensation language. Examine any individual instance of the use of a sensation term, and an account of such a use given exclusively in terms of teaching links – or at any rate in terms of teaching links with some qualitatively unspecified 'sensation' thrown in for good measure, may seem plausible. People appear to be reporting on observable states of affairs when they say, watching other people, what. mental state they are in. People reporting on their own experiences may, if closely questioned, find themselves unable to characterise these experiences without having recourse to externalities. All this is of course to be expected according to **[54]** C-crude just as much as according to C-subtle, which is why it is so difficult to insist on C-crude, the more philosophically daring of the two theories. The crucial assumption that underpins sensation talk just does not reveal itself on everyday occasions when sensation talk occurs. The form of life which can apparently be adequately described instance by instance in terms of teaching links is, however, powered by a faith that these links are indeed *links*, a faith in that which they are links *with* – not only in its existence, but in its marching qualitatively in step with the links. I have conceded for the sake of argument that this powering might be more properly described as psychological than as logical. But having made this concession, I maintained that the importance of the assumption was not affected. This I can now elaborate. Concepts need to be useful as well as coherent: they are kept alive by our need to use them for certain ends. And although they might, logically, equally well subserve other ends, these may not be ends which appeal to us. This being so, it is perhaps immaterial, if a concept dies, whether it dies a logical or a psychological death. A couple of examples may help to clarify this distinction.

If it is shown that the concept of free will is incoherent; that it appears to point to a genuine possibility which, however, evaporates as soon as it is subjected to close scrutiny; then that concept dies a *logical* death. It is not merely that we cease to believe (say, because we come to accept determinism, and take it to be incompatible with free will) that free will exists, but rather we realise that we simply don't know what free will **[55]** *would amount to*. So the notion can no longer have any intelligible use. The death that would be died by sensation concepts, however, if people came to believe that C-subtle had to be the correct theory of their meaning, would be *psychological*, in this sense: it would still be *coherent* to continue using the concepts. Although their subjective foundations would be largely washed away, the *publicly enforceable*[41] rules for their use would be unaffected. To this extent they would still be the same concepts, and if they fell out of use, it would not be because their use had become *senseless*, but because it had become *pointless*.

That a language should serve some interest or other is in practice as vital to it as that its concepts should be coherent: so we might well, subsume both kinds of conceptual obsolescence under one heading, even though they take two different courses, and are to that extent of two different sorts. Compare in this connection something said by Hacker in interpreting Wittgenstein (p. 241): 'it is a deep misconception to think that one can account for the sense of a sentence independently of its point. This is like thinking that one can give an adequate description of chess without involving the notions of winning and losing'.

Although in particular practical cases where decisions are required it may sometimes be fair to say, in C-subtle mood, that philosophical-sounding questions about knowledge of other minds show, if not mere perversity, ignorance of the language, this does not affect the fact that behind the everyday use of this department

[41] The *private* rule of rigid designation *wouldn't* be unaffected: but this rule is held by C-subtle to be senseless, so the conceptual murderer, from his own point of view, is elucidating an actual concept rather than providing us with a new one.

of language lies an assumption about the correlation between **[56]** sensations and teaching links. It makes sense both to assert and to question this assumption; and without it, many questions in this area would lose their point. For example, many moral decisions depend on assessments of degrees of suffering and pleasure caused by different possible courses of action. Do those who are already well acquainted with hardship suffer less from new hardships than those who have previously had a comfortable time? Would the pleasure caused by a given action outweigh the suffering the action would also cause? And so forth. If suffering and pleasure were reduced in a Rylean sort of way to behavioural patterns and dispositions (and in effect they would be, if C-subtle were carried through), the whole point of avoiding the former and pursuing the latter would be lost. Somewhere along the line the buck must stop, and moral behaviour be grounded in the desirability or otherwise of experiences. Moore was right to say that, in the last analysis, only states of mind matter.

It should now be clear why it makes little difference whether we call the assumption in question a psychological precondition or a logical requirement. For simplicity's sake (to put it at its lowest) I shall stick to the course chosen in my initial specification of C-crude-subtle, and call it a logical requirement.

C-crude-subtle, then, is C-crude refined in such a way that it fits better with our actual linguistic behaviour, particularly as far as the role of teaching links is concerned, without losing the particular connections between sensation terms and qualitatively distinguished kinds of sensations, connections which are the lifeblood of sensation language. With these refinements made, and given (cf. chapter 1) that C is a possible theory, I suggest that C-crude-subtle provides the basis for the true analysis of our sensation language. **[57]**

The Ambiguity of C-subtle

According to C-crude-subtle, as I have said, teaching links are the means of teaching the use of sensation terms, because they fix the reference of those terms, but they are not definitive of the sensations whose names they help to teach. One reason why C-subtle is so persistently seductive is that it can be taken as ambiguous between C-crude-subtle and the theory it openly claims to be. 'Sensation of a type which has such and such teaching links' can be taken in two ways, depending on whether the 'which has' is taken as reporting a contingent or a logical connection between links and sensation type. Either the links are definitive of the sensation, whose subjective nature is immaterial, or the sensation is defined by its subjective nature, and the links are the vital means by which, thanks to the bounty of nature, we are empowered to learn which terms go with which subjective referents. While with one's mouth one espouses the first of these two interpretations, in one's heart one treasures the second, and feels secure: this is the temptation for a C-subtle supporter. If the full ramifications of the first interpretation were faced up to, then the theory would not be so calmly adopted.

It is plain that the proponents of C-subtle do not really believe in the first interpretation of their theory except when immediately concerned to defend, and so conscious of, its counterintuitive content. I shall now examine two passages, one from Pears and one from Locke, which show (among innumerable others) that deep down they are more sensible than they would like us to believe.

Here is the passage from Pears (pp. 154–5):

[58] What fills [the promiscuously open, but never really embracing arms of the variable 'sensation of a type which …'] in each person's case is a matter of indifference to the meaning of the phrase. All that C-subtle requires is that, if, unknown to us, sensations of pain did vary from person to person, at least they would have to remain constant for each particular person. For if they suddenly and completely changed their

character for a particular person, he would become a baffling case, because he could always with perfect sincerity deny that he was in pain, in spite of the fact that all the public criteria for his being in pain, except, of course, his own verbal reactions, were fulfilled.

This 'baffling case' is closely akin to the case I considered above (pp. **22**–**23**), where a man who has never felt pain before feels his first pain in the absence of any of its normal teaching links. To Pears these cases indicate that a special requirement must be imported into C-subtle which, I shall argue, has no place there. To me they indicate that it is just as well that sensations and links do normally stay paired, for if they didn't, cases of this kind would indeed be baffling to a language community which takes for granted that links of the right kind are always at hand. Since the workings of our concepts are dependent on the world staying as it is in this regard, I am thankful that the world does stay as it is: Pears, noting the same dependence, wishes to re-doctor for the intrasubjective case the already doctored concepts of C-subtle so that they would be immune to the radical changes envisaged in the cases considered. But this is only a procrastinating manoeuvre: at some stage the dependence of our conceptual machinery on the nature of the world will have to be acknowledged, so why not acknowledge it at the outset, in the first layer of the onion? But, such considerations aside, Pears's manoeuvre can be disallowed on internal grounds. To a serious, consistent **[59]** adherent of C-subtle the requirement Pears here wishes to import is not available. In the first place it is a mere stipulation, no justification being offered for why we should believe it to be true, and in the second place it doesn't even make sense, as far as C-subtle is concerned. The talk of the 'character' of a sensation in the quotation from Pears is not C-subtle talk, but C-crude talk. According to C-subtle, the character of a sensation is given entirely by its teaching links. In fact 'character' *means* 'teaching links'. It is simply meaningless to talk of some other sort of concealed 'character' which a sensation has, for it is the burden of C-subtle's argument that such a 'character' could not be discussed, there being no teaching links to

endow the term with meaning.[42] If we take this implication seriously, a lot of what. has been said in this thesis becomes nonsense in the same way: wherever the 'quality' or 'subjective nature' of a sensation is mentioned, even if only to say that it plays no part in the meaning of a sensation term, the impression that we understand what is being said must be a mere illusion – or else the terms in question must be analysed in terms of outward phenomena, as C-subtle requires, in which case something utterly different is being said from what was intended. For example, read in this light the first quotation **[60]** from Pears on p. **15** is nonsense: 'different kinds of sensation' can't mean what Pears[43] needs it to mean, viz. 'subjectively different kinds of sensation' – it can only mean 'sensations with different kinds of teaching links', for the term 'sensation' is an empty variable, with no sense beyond that given to it by its teaching links. But if this has to be the meaning of 'different kinds of sensation', Pears is saying 'if […] the same teaching links pointed to sensations with different kinds of teaching links in different people': and this is self-contradictory. So he cannot make the point he is trying to make, that 'those differences would not be picked up in the meaning of the word "pain".' All the differences C-subtle can specify *would* be so picked up.

Proponents of C-subtle cannot have it both ways. Once they have disallowed a connection between a particular subjectively

[42] It would be possible to espouse C-subtle and still hold that talk of the subjective character of sensations was meaningful. One might agree that sensations did have such a character, and that this character might be intersubjectively similar or different, but hold that there was no way of knowing, and that in any case the logic of sensation terms was such that it didn't matter. Such a moderate version of C-subtle may seem more plausible (though it is still opposed to the intuitions I am defending). But I am treating C-subtle as essentially deriving from that philosophical stance which regards talk of the subjective as meaningless, and is thus forced into C-subtle (if not right into behaviourism), rather than taking it up as an option on an equal footing with C-crude-subtle. When motivated in this way, C-subtle is committed to the absurdities I discuss in the main body of the text.

[43] Here and *passim* hereabouts, cf. footnote to p. **16**.

differentiated type of experience and the meaning of a sensation term, all talk of sensations varying from person to person or within a person is nonsensical. If a sensation is defined in terms of its links, the only residual difference possible between one person and another that can be meaningfully spoken of is that one person may have the links *and* the sensation, and another the 'links' but no sensation. To say that someone has a sensation but no links,[44] or that a sensation may be unfaithful to its links, makes no sense. **[61]**

But of course we (Pears among us) do not want nonsense to expand its domains in this way. *Of course* Pears's requirement that sensations of pain remain constant for each particular person makes sense (even though he does not justify it in terms of C-subtle); *of course* we understand all the talk about subjective experience perfectly well; *of course* the first quotation from Pears on p. **15** is intelligible and says what he wants it to say. So the theory which commits us to yielding before such nonsensical imperialism must be false. And it is. Here is the passage from Locke (p. 97):

> Whether or not it is possible to tell whether pains other than our own exist, the fact is that there are public ways of telling when something, given that it exists, is a pain. How do I know that this, which I feel, is a pain? Because it feels like one. How do I know that sensations which feel like this are pains? Because they are sensations of the sort I get when I am wounded, which make me want to cry out, etc.

Compare here the quote from Locke on p. **17**, where he says, for example, 'the sort in question is determined not by how it feels', and 'a sensation's being a pain sensation is not a matter of how it feels'.

[44] We cannot give a sense to such a remark by analysing 'sensation without links' as 'that which there normally is over and above such and such publicly observable phenomena, but in this case without the phenomena'. For this would again be to break the rules of C-subtle, by assuming that 'that which there normally is' has any sense: the character of 'that which' is entirely parasitic on the character of the links, so it cannot have any 'normal' character of its own which would serve to identify it in the absence of the links. Its character is always the same and never detachable.

Locke has given himself away in the same manner: 'it feels like one', if C-subtle is consistently held, must mean either nothing, or else (a *reductio ad absurdum* of C-subtle, surely) 'it has the teaching links definitive of pain'. But Locke is not taking it like this, or his next two sentences would be superfluous and tautologous. He is taking it as I take it, as meaning that the subjective nature of the experience shows that this is a case of pain. It is made easier for him to perpetrate this sort of inconsistency by the ambiguity in formulations of C-subtle. Locke's 'sensations of the sort I get when' in the above quote is ambiguous in just the way specified earlier. When he needs to **[62]** be sensible, it can be understood as 'that type of subjectively identifiable experience which as it happens I get when', and when he wants to defend C-subtle it can be understood as 'that class of sensations whose members are defined as those which I have just when'. If the ambiguity is rooted out, C-subtle is seen in all its bleakness.

C-subtle-intermediate

It may be argued that I have neglected an intermediate interpretation of C-subtle which is in fact the one espoused by Pears and Locke. This interpretation would have something in common with the first of the two interpretations already given (strict C-subtle), something with the second (C-crude-subtle). It would avoid the emptiness of the first, and the vulnerability of the second. It would be the theory that the term '*x*' must, for any particular person, always refer to the same subjective experience, though *which* subjective experience is determined not by the subjective nature of that experience, but by the links in whose company it is normally found. Once, however, it has been established, by consulting the links, which subjective experience is at issue, reoccurrences of the sensation can be subjectively identified: it becomes definitive (though parasitically on the links) of the sensation that it should feel a certain way. Hence Locke's 'it feels like one' is given a sense.

On this view it is of no significance for the meaning of '*x*' if the links in my case are regularly accompanied by an experience subjectively different from the experience which accompanies them in your case: such a difference would be 'left on the plate untouched' (Pears, p. 152). But it *does* matter that in my case the sensation which accompanies the links should be subjectively the same or similar **[63]** on different occasions. Any differences which arose here would give rise to 'baffling cases', and would indeed affect the meaning of '*x*', which is 'that subjective type of sensation which, for any particular person, is accompanied by such and such teaching links'. The requirement that *within a given person* the subjective nature of '*x*' should be constant is what this version of C-subtle has in common with C-crude-subtle. The lack of a requirement that there should be such constancy *between different people* is what it has in common with strict C-subtle. The former requirement gives it a respectable foothold in the inner world, the lack of the latter a chance of avoiding the standard objections to C-crude. Such a theory I call 'C-subtle-intermediate'.

I do not deny that such an interpretation of C-subtle may be what Pears and Locke have in mind.[45] But whether these authors

[45] If it is part of the meaning of sensation terms that for any given person a given sensation must be subjectively similar on different occasions, then if the links–quale pairing changes for a person, and if he learnt the meaning of a sensation term before the change, he is going to be in an impossible position. This is either an illustration of how important it is that the regularity of nature doesn't subside; or (if theories are to be designed to cope with the maximum number of possible worlds) an argument in favour of C-crude-subtle or C-subtle as against C-subtle-intermediate.

If the links–quale pairing according to which we learnt our sensation language as children changes, either we continue calling the same qualia by the same names (C-crude-subtle), or we continue making the links definitive of sensation terms (C-subtle), but we cannot do both (C-subtle-intermediate). We could take this possibility not as an argument against C-subtle-intermediate, but as a further illustration of the point that, for all forms of language, it is possible to imagine situations in which we would just not know what to say. But given that there is a theory (C-crude-subtle) which would enable us to know what to

[64] intend it or not, it will not do, for reasons which should by now be obvious.[46] Such an intermediate theory wants to have its cake and eat it. For arguments which make it impossible to require that the subjective nature of *x* shall be the same between different people also make it impossible to require that it shall be constant within any particular person: in neither case are there adequate public criteria available. Conversely, the moment one allows the possibility of requiring that the subjective nature of *x* be constant for any particular person, it is equally possible to require that it be the same between persons (for further argument to this effect see 'Intrasubjectivity and intersubjectivity', the last section of this chapter, below). Any theory which adopts one requirement and rejects the other is inconsistent, *unless* it admits that the requirement it rejects is perfectly possible, but undesirable on other grounds. But since C-subtle rejects the between-people requirement precisely on the grounds that it is an impossible requirement to enforce (and so senseless), its proponents cannot escape through this loophole. Either the other requirement is impossible too; or both are possible. The latter is the case: both requirements can be made; should be made; and of course are in fact made.

Intrasubjectivity and intersubjectivity

The disanalogy between the intrasubjective and the intersubjective case is sometimes believed to be sufficient to make it impossible to give a sense to the expression 'same sensation as', conceived according to C-crude, intersubjectively, even if the arguments designed to establish that the expression is intrasubjectively

say if this possibility should occur, and which is the best theory on independent **[64]** grounds, we do not need to be resigned to the possibility of such bafflement. Of course, even according to C-crude-subtle, should the links–quale pairings start going berserk, then teaching new language pupils sensation language would be a practical impossibility (without neurophysiological help). But as I have already said, there comes a point when it has to be admitted that the linguistic status quo is dependent upon the status quo of nature.

[46] [Expand these reasons: they are by no means convincing.]

meaningful are successful. For example, one can appeal to memory in the case of within-subject comparisons, but not in the case of **[65]** 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. If these considerations are accepted, it will not be possible to attack C-subtle-intermediate on the grounds of straight inconsistency, as I have just done. For to such an attack it will be replied that the conceivability of requiring that the subjective nature of *x* be constant for any particular person is not sufficient to guarantee the conceivability of making a similar interpersonal requirement. Naturally, since C-subtle-intermediate is identical with C-subtle as far as the interpersonal case is concerned, I hold that it too is vulnerable[47] to the argument I have already advanced against C-subtle in chapter 2 (but cf. note to p. **51**): C-subtle does not achieve immunity by drawing up a compromise agreement with C-crude-subtle. But even so, to refute C-subtle and C-subtle- intermediate is not to establish C-crude-subtle. And C-crude-subtle must reckon with the considerations just mentioned.[48]

The crucial argument is that it does not make sense to talk of comparing two subjective experiences which are necessarily experiences of different subjects. If I logically cannot compare your experience of *x* with mine, either by having the two together, or by remembering a past instance of one and comparing the memory with a current instance of the other, what is the cash value of the question (or of either of its answers) whether the two experiences are subjectively comparable?

First, even if the premise about the logical impossibility of having someone else's sensations is true, I am not moved by this argument.[49] I do not think that logical privacy delimits the field **[66]** of coherent comparison in the way suggested. The fact that

[47] [*Is* it? If so, make this point on p. **55**.]
[48] [This point should appear as less of an afterthought.]
[49] [Dubious case made in what follows.]

what is mine is mine and what is yours is yours is too trivial to be such a barrier. For example, although my face is mine, and yours is yours, we can still perfectly coherently compare the two faces in respect of all their properties. And it should be possible at least to *understand* the notion of comparing my sensation and yours in respect of their subjective quality. The fact that we cannot in practice effect a direct comparison does not make the idea of comparison nonsensical, any more than it would be nonsense to speculate about similarities between my face and that of Alexander the Great – even though, he being dead, it is impossible actually to effect any direct comparison. Your experience cannot *be* mine perhaps: but why should that prevent it being *like* mine?

But if this idea still sticks in the throat, it is possible to go further and argue that the possibility 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. The contingency in question is universal, perhaps; but it is still a contingency, though its universality is precisely the condition under which it is most likely to be mistaken for a logical necessity.

The requirement that one occurrence of *x* should be subjectively like another can be validated in the intrasubjective case by drawing attention to psychological continuity, 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 be tween the two cases seems to depend on the fact that psychological continuity occurs only intrapersonally. If this apparently **[67]** 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 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

possible cases which our present notion of personal identity would simply be inadequate to describe. So 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 halves 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. And yet such continuity has always been one of the main criteria of personal identity. This is possible as things are, but it is wrong to think that it is universally satisfactory: for psychological continuity is not logically indivisible, as personal identity is – it can branch, as this example illustrates. 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 comparisons are possible, it makes sense to say of both the resulting **[68]** 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 a sense to 'same sensation as' intersubjectively. 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 neither subject has any prehistory in common with the other. For these cases we need to imagine 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.

Does not this achieve our aim? I think it does. It is a way in which I could compare your experiences and mine. The possible objection is this: how do I know that post-fusion experiences are comparable with either set of pre-fusion experiences? Suppose the two pre-fusion experiences of *x* were in fact different: what would the post-fusion experience of *x* be like? This rather depends on what sort of experience *x* is, and 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 of *x* are incompatible: so the post-fusion experience of *x* will be a compromise between the two – 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. **[69]**

Imagine a different system. In fusion, half of each pre-fusion brain is put in a new body (or one of the old ones, or a combination: 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; sensation representation is contralateral (the left side of the body being represented in the right hemisphere, and vice versa); the quality of sensation in the unitary post-fusion centre of consciousness is determined by the structure of the hemisphere in which that sensation is represented; the two pre-fusion experiences of *x* were qualitatively different. Here we have a case where direct comparison of bodily sensations will be possible: if I put both feet in hot water, the sensation of heat may be different for each foot.

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

are not subjectively contradictory,[50] then 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. If they are contradictory, he will find himself unfeeling in that regard, and thus know at least that the old experiences were different, if not in what way. I have dealt with these examples sketchily. They merit a whole book. But I hope I have said enough to suggest that the **[70]** 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.

Finally I want to draw attention to some things that Strawson says in the chapter about persons in his book *Individuals* (1959). He adduces some logical considerations which count against the argument that possible intrasubjective requirements are not necessarily possible intersubjective requirements: these considerations tend to show that C-subtle-intermediate is logically impossible. If the best weapon to fight logic is more logic, then here is a useful weapon.[51]

Strawson's arguments are complex, and I do not attempt to represent them self-sufficiently here. I just draw attention to his conclusions, and note their implications for the intrasubjective–intersubjective dispute.

The crucial 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 to 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: and no attenuated version of the concept is

[50] [Could they be?]
[51] [Cf. p. **65**.]

workable. From this it follows that P-predicates (predicates peculiarly applicable to persons, including sensation predicates) essentially have first- as well as second- and third-**[71]**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 ascribe them *in the same sense* (fig. 2, p. **62**).

But, as Strawson notes, and as C-subtle-intermediate accepts, there are some P-predicates, sensation predicates among them, which I ascribe *to myself* on the basis of non-behavioural evidence ('it feels like it'), 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 now look carefully at just what this requirement of conceptual parallelism involves, we shall see that it is fatal to C-subtle-intermediate, for it commits us to that very intersubjective rule of phenomenal quality which C-subtle-intermediate disowns, as follows.

In one's own case it is required by C-subtle-intermediate that occurrences of sensation *x* should all be subjectively similar. So far, so good: C-subtle-intermediate 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-intermediate 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. It is 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 others 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 **[73]** will be the same as it is for me. For although they will need to have the same subjective experience on all occasions of the occurrence of *x*, it need not be subjectively similar to the experience *I* always have when I have *x*. One of the

necessary conditions for the applicability of '*x*' will be different in the two cases and Strawson's requirement is not met. **[72]**

Fig. 2 *Strawson's view of P-predicates*

[73] An objection has been building up during this last paragraph, and must now be allowed to burst forth. It will be said that I have confused C-subtle-intermediate with C-crude-subtle; that the extra requirement I have just discussed is no part of C-subtle-

intermediate, and I am not entitled to include it as one of the necessary conditions for the applicability of '*x*', so long as I am discussing that theory. There are as many different ways of feeling the same as there are different sensation qualities, and the requirement that C-subtle-intermediate *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-intermediate does pass the Strawsonian test.[52]

Why do I not accept this objection?[53] Because in my own case, once I have learnt via the appropriate teaching links which subjective kinds of sensation to pick out as the proper bearer of which names, then the way a sensation feels becomes a criterion **[74]** of its identity.[54] Once *this* sensation 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 just *is* a necessary condition for the applicability of '*x*' in my own case. 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. But this is arbitrarily to exclude one

[52] [This may be right.]

[53] [What follows is sophistical?]

[54] I say '*a* criterion', not 'the criterion', because for C-subtle-intermediate teaching links are also criterial. These two criteria will not inevitably always give the same answer, and this puts C-subtle-intermediate in an irresolvable dilemma: as I have pointed out, if the links–quale pairing were to change in the case of a person who had already learnt the use of sensation terms, he could not satisfy *both* the requirement that all occurrences of a sensation should be similar-for-him, *and* the requirement that which sensation should be given which name is to be determined by the teaching links. But the fact that these criteria would conflict in what I hope is a hypothetical case does not mean they cannot both be criteria (albeit over-determining criteria) in actual cases.

criterion for the applicability of a term from playing a role in that term's meaning.[55] Such a move is akin to the attempts to separate meaning and reference which were discussed and rejected in chapter 2.

There are indeed expressions which refer to particulars under descriptions which do not allude to the particulars' particular qualities. To be Queen of England one does not have to be the same person as some particular Queen of England. But although being Queen of England is not a matter of being a particular person, it is my thesis that being a certain sensation *is* a matter of having a particular subjective quality. We sometimes need to **[75]** refer to Elizabeth II qua Queen, or qua the monarch: but we do not want, in general, to refer to particular kinds of subjective experience simply qua the associates of certain teaching links.[56] We want to refer to them qua having the subjective qualities they have.

I don't mean to imply by this that C-subtle-intermediate is possible though ill adapted to our interests.[57] strict C-subtle, I have allowed, is a logically possible theory:[58] but C-subtle-intermediate is not, for the reason given **three paragraphs previously**. One cannot prise apart *feeling the same* and *feeling like this* in the same way that one can prise apart *being this particular person* and *fulfilling this role* (or whatever the criteria for being the English monarch are). At least one cannot, separate them in a list of the criteria used by me to establish that I am having a particular sensation – criteria which must also be used by others if Strawson's symmetry requirement is to be met. Once we let in the intrasubjective criterion that the quality of x should be constant, this commits us to the intrasubjective criterion that the quality of x should be like *this*; and this in turn commits us to the intersubjective requirement that the quality of x should be comparable, i.e. like this, for all persons.

[55] [Is it? This may be a weak point.]
[56] [C-subtle(-intermediate) would say we do.]
[57] [But it is, perhaps.]
[58] [Amplify. It is possible only if it tempers its claims (cf. p. **29**, p. **51** note).]

Strawson's logical arguments are directly opposed to the argument from logical privacy considered earlier. Since the latter argument has already been faulted on other grounds, it is not difficult to decide which competitor should go to the wall. **[76]**

APPENDIX TO CHAPTER 3

The four faces of C

It would not be hard to confuse the various versions of C which have been discussed thus far. So it may be worth clarifying what they amount to. This I shall do in two ways. First I shall show diagrammatically what they would have to say about three important stock hypothetical cases which crop up repeatedly in discussions of the mechanics of sensation language. Then I shall give a summary of the specification of each of the versions of C. The summaries may be fleshed out by comparing them with the diagram, so as to see what goods they deliver in concrete situations. There are four versions of C in question: two versions of C-crude (C-crude-crude, the one quoted from Pears at the beginning of chapter 1; and C-crude-subtle, my refinement of this), and two versions of C-subtle (strict C-subtle and C-subtle- intermediate). **[77]**

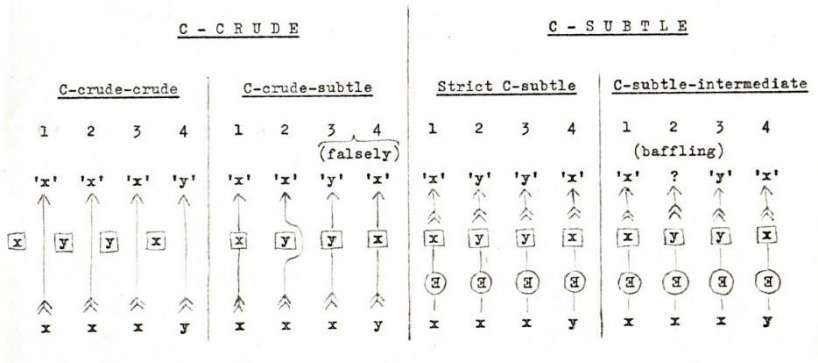

Fig. 3 The four faces of C

[76] Let me explain the symbols and labels used in the diagram.

The numbers in the first row are the numbers of the cases: one normal case and three *lusus naturae*. Cases 1, 3 and 4 are all learning situations. Case 2, as specified, is a situation where the term has already been learnt, and is now in use: but nature is playing a trick. These cases are as follows (so as not to beg any questions, I describe the cases in terms which assume that C-crude-subtle is true: indeed, if strict C-subtle were true, it would be impossible to describe certain aspects of them):

1. Person A in a paradigmatic sensation term learning situation. There are teaching links available, and they are the right ones to go with the accompanying subjective experience, *x*.

2. Person A again, and sensation *x* again, after the meaning of '*x*' has been successfully learnt. But here nature has gone berserk: teaching links are present, but not the right ones for *x;* they are the ones that normally go with y.

3. Person B with the same subjective experience as A, but different teaching links regularly present from the start.

4. Person C with the same links as A in case 1, viz. those appropriate for *x*, but a different accompanying subjective experience, viz. *y*.

The next row in the diagram is a row of names: it shows which name each of the four versions of C would be obliged to give to the sensation in each of the four cases described.

The third row shows which links are present. The square boxes just serve to distinguish links from subjective experiences, and from names (which are in turn distinguished by quotes). The letter inside the box shows which sensations the links normally accompany.

The fourth row shows which subjective experience is present. Letters are here used as rigid designators of subjectively differentiated experiences. **[78]**

The symbol ↠ means 'defines as such'. E.g. '*x*↠"*x*" ' means '*x* is called "*x*" because it is *x*.'

'x–(∃)–[x]' means 'the necessary condition that some sensation exist in order that the presence of potential teaching links constitute a paradigmatic learning situation is satisfied for teaching links [x] by sensation x.' The relation –(∃)– must, in the case of C-subtle, obtain between a sensation and a set of links before the relation ↠ may be set up between the links and a sensation term.

'x↠[x]↠'x' means 'links [x] act as teaching links for learning the defining connection between x and 'x'.

The diagram shows that C-crude-subtle is the one theory of the four which, if nature plays up, leads to misidentification and so misnaming of sensations according to its own rules. (Both versions of C-subtle misname sensations according to the rules of C-crude: but as far as they themselves are concerned they name correctly in all four cases, excepting C-subtle-intermediate's hesitation over case 2.) But I accept this consequence, regarding it as a recommendation of C-crude-subtle that it alone makes provision for making permanent mistakes[59] in identifying subjective experiences. This is just the feature that makes the theory unacceptable to those philosophers who adopt C-subtle, or reject C altogether, for they reject on principle theories which make it possible for us to be undiscoverably mistaken.

Summary of the four theories

C-crude-crude says that the subjective nature of a sensation decides what name it shall have, and that the name is given to the sensation in a private naming ceremony, without the mediation of teaching links. Rather, teaching links are not even mentioned by this theory, even to be denied a role. But I have put them in the diagram, unconnected, to show that they have been neglected.

Strict C-subtle says that, given only that some sensation exists, it is given the name it is given solely in virtue of the teaching links in whose company it is found. The links will determine the name

[59] It is always possible to make readily corrigible mistakes through carelessness, or ignorance of the links criterion.

of the sensation in every individual case of the occurrence of a sensation, for there is no sense in the supposition that, once a sensation term's meaning has been learnt, it might afterwards be applied in virtue of the subjective nature of the sensation. This is because the notion of a subjective nature is empty: the only nature a sensation may intelligibly be said to have is that bestowed on it by its links.

C-subtle-intermediate says that, for any particular person, which of his kinds of subjective experience is chosen as the bearer of a given sensation name is determined solely by the teaching links which regularly (d.v.) accompany it. But once this relation has been set up, then the name must continue to apply to the same **[79]** subjective experience throughout that person's life: accordingly it will not be out of place, after the term has been learnt, to apply it in virtue of the presence of the appropriate sensation, disregarding the links which nevertheless, no doubt, are faithfully dogging its heels (but cf. case 2, where there seems to be an irresolvable conflict of criteria).

C-crude-subtle says that the subjective nature of a sensation is what. determines its name; that this nature must remain the same within a person's sensational history; that it must be the same in different people. But it also recognises the fact that teaching links are necessary in order that the meaning of terms conceived on this model may be learnt: note however that these teaching links will only perform their function successfully if they are correctly paired with the right kind of sensation to bear the name they point to. We depend on the regularity of nature for the correct use of sensation terms.

C-crude is unworkable; strict C-subtle has little to do with sensations; C-subtle-intermediate is neither consistent nor logically possible; C-crude-subtle is the truth. **[80]**

4 TEACHING LINKS

As I have already said, I have been operating so far, for the sake of making the issues as clear as possible, with a somewhat idealised notion of sensation language. There are very few sensation terms – perhaps none – whose meaning is entirely dominated by the subjective nature of their referent. This shows up in the fact that it is difficult to think of a list of sensation words more than two or three words long, unless one is allowed to include those sensations that can be described only by borrowing the name of a publicly observable object, event or phenomenon, in some way closely associated with them, either by being their cause, or their effect, or their expression, or by having some less direct connection with them which can nevertheless be used to underpin the use of the name of the public item as a label for 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 used to refer to the 'pure' sensation by itself, neat, alone. This may be 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 the same sort of sensation occurs all over the body, so that we could not use as a general name for it a term drawn from one particular bodily area: but this is a complicated issue, and I do not pursue it here.

It may be thought that my concession that almost all sensation terms have meanings in which teaching links feature to some degree or other amounts to an abandonment of C-crude-subtle, at any rate in an unadulterated form; and a compromise with C-subtle. But this is not so. And it is time to show that it is not so, by **[81]** considering sensation language in a less idealised form: by looking at teaching links more closely, and discussing their role

more realistically. I have made too much use of a vaguely characterised, supposedly intuitively clear notion of 'teaching link', and it is time to make the notion more precise, so that it can be shown how C-crude-subtle can be reconciled with the involvement of teaching links in the meaning of sensation terms. To this end I want to ask three questions about teaching links, in order to clarify respectively their nature and function, their connection with meaning, and their relationship to the referents whose names they make it possible to teach.

The three questions are:

(i) What *are* the teaching links of sensations? (There are undoubtedly different kinds of links, and the same things cannot necessarily be said about all of them.)

(ii) What is the connection between teaching links and the meaning of the terms taught with their aid?

(iii) What, in the light of the answer to (ii), is the importance for the meaning of a sensation term of its private referent?

I take them in order.

(i) *What are the teaching links of sensations?*

> It is necessary to emphasise the variety of teaching links.
> Pears, p. 149

First, some names of sensations, starting with the 'pure' ones and finishing with the most links-ridden ones: a sketch of the field of operations of teaching links.

Pain, itches, aches, agony, discomfort, giddiness, cramp, tinglings, throbbings, nausea, fatigue, glowings, tension, lethargy, palpitations, shivering, hypersensitivity, shock, taste, **[82]** sound, pins and needles, colour sensations, sensations of cold, of heat, of being tickled, of wanting to yawn or laugh, of being sick, of illness, 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, of sweating, of wind on the flesh, of orgasm, of indigestion, of a full bladder, of hunger, of thirst, of being emotionally disturbed, electric shocks, stomach ache, the

'jumping' sensation on falling asleep, sensations of fear, of anger,[60] and so on ad infinitum, or nearly so. This is a very mixed bag, and it may be argued that some of the earlier names are not really names of sensations at all: for example, cramp may be thought to name a physiological condition, so that I should have said 'sensations of cramp'. But sensations can borrow names from their causes – 'fatigue' is another example – and in any case nothing hangs on the omission of 'sensations of' in some cases. It may be added if it is felt to be necessary.

In most cases the function of the links, and what the links are, is self-evident: the name of the sensation is a giveaway. But in other cases the identity and role of the links need to be spelled out, for there is no superficial verbal evidence to hand. I shall discuss one or two cases, without pretending to be comprehensive. In particular I shall not say anything here about non-bodily 'sensations' such as visual sensations: these are special cases which require special treatment – it may even be that they are inappropriately called 'sensations'. I will mention them in chapter 5, but only cursorily.

Thanks to Wittgenstein, the bulk of discussion of teaching links takes place with reference to pain: and, also thanks to Wittgenstein, the best-aired teaching link of pain is its natural **[83]** expression. Pain is as good a sensation, and pain's natural expression as good a teaching link, to start with as any other when specifying the nature of teaching links for sensations in general, as long as one does not get stuck at this point of departure. Pain has other links, and other sensations have links of different kinds from the ones pain has.

With this caveat in mind, I too shall start with the natural expression of pain. When people are in pain they tend, if the pain is severe or their tempers frayed, to 'wince, scream, flinch, sob, grit their teeth, clench their fists, exhibit beads of sweat' (Putnam (1965), p. 9) and so forth. Assuming that these responses are innate

[60] [Emotions are differentiated more by their intentional content than by their phenomenal nature.]

(and surely at least some of them are), we can say that they provide natural links for teaching the use of 'pain' (on any theory of the meaning of 'pain' except for C-crude-crude: I am not *here* concerned to adjudicate between them). I hesitate slightly to call beads of sweat an 'expression' of pain: rather it seems to be an effect. How is one to distinguish between expression and effect? It is not immediately clear, but the distinction appears to be connected with that between actions and passive reactions. Wincing is something one does:[61] sweating is something that happens to one. But the distinction is not important for the language pupil.

As Pears points out, there are refinements in pain talk which can't 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): these descriptions can be subdivided into two groups. The first group describes the pain **[84]** by analogy, by implying that it has a quality like the pain one gets when one is really being burnt, prickled, chilled etc. The second group draws attention to the pain's pattern: a throbbing pain has a certain time pattern of intensity; a shooting pain travels a fair distance in the body at high speed; a stabbing pain is sharp (another analogical description) 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 narrow and definite location, and intensity? What sort of analogy (if any) is 'dull'? And it is sometimes easy to make a mistake: 'stabbing' can be taken 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?

[61] [Really? At most in part.]

Strangely, when Pears talks (p. 148) of the sort of distinction which can be taught with reference to the 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 surface and interior, and the distinction between interior 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.)

Further distinctions rely on the fact that certain sensations occur in company with certain public events (either before or after them), but are yet neither their causes, nor their effects, nor expressed by them. There are many examples of sensations picked out in this way in the list on p. **71** above. Pears mentions **[85]** the sensation that normally immediately precedes a yawn: it is not the cause of the yawn, nor is the yawn its natural expression, but it can be safely characterised as a yawn's predecessor. A similar case is the distinction between pain and nausea – nausea is an unpleasant sensation which may produce all the 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 constant harbinger of vomiting, an impeccably public event. Also in this class there is that sudden muscular jump one's body often gives as one is dropping off to sleep: unless this experience occurred regularly at this or at some other fixed point in our personal histories, it would be much harder to communicate about it without being in doubt as to whether we were all talking about the same sensation. It would not be sufficient to characterise it as the sensation of a muscular spasm: there are spasms and spasms, and this one feels special.

Finally one should mention 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 analogical descriptions of sensations like 'pins and needles': 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 intensity from one field of sensations to another. The intensity of a 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 an intense sensation than there is of a mild one. But just as we can readily transfer the concept 'more' to subject matters widely different from those in connection with **[86]** 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.

Such in 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 expressions of pain led him, as Pears says, to treat 'what are really descriptions of sensations as if they were verbal expressions of sensations'. Perhaps we are by now sufficiently aware of the multiplicity of links not to make any more mistakes of *this* kind at any rate.

(ii) *What is the connection between teaching links and meaning?*

> From the fact that certain contingent circumstances have in general to obtain before a certain word can come into use (even if this fact is itself logically demonstrable) it by no means follows that these same contingent circumstances are logical conditions of the word's correct use on a particular occasion, or determine what its meaning is on that or other occasions. Thus 'pain' might be the name of a completely private experience (a word which could be legitimately used whatever was happening overtly) even though the word could not have come into use unless, normally, these private experiences were correlated with overt occurrences. The point is of some methodological importance.
>
> Hare (1972), p. 89

To what. degree are the teaching links by means of which a word is taught logically implicated, or liable to become logically

implicated, in the word's meaning? Here we shall have to consider the problem of the 'criterial' relation which is discovered in Wittgenstein; and the question of the relationship between teaching **[87]** links and 'criteria'. So this section is divided into two. First I shall ask what sorts of connections there can be between circumstances which justify the use of a word and the meaning of that word. Then I shall ask which of these sorts of connection seems to hold between teaching links and the sensation words they help to teach.

Criteria and meaning

We find in the literature a variety of terms used to describe the circumstances which prompt the use of a word: the critical distinction, in terms of which the use of the other terms can be specified, is that between 'symptoms' and 'criteria'; and it is on this distinction that I shall concentrate. I use 'criterion' in a somewhat technical sense: of course in common parlance 'criterion' could cover symptoms as well, but in the literature it is used to specify circumstances which supposedly stand in a special logical relation, whose supposed nature will shortly emerge, with the term whose use they prompt. 67

The relation between symptoms and the term which refers to that of which they are symptoms is comparatively straightforward and uncontentious. Wittgenstein writes: 'I call "symptom" a phenomenon of which experience has taught us that it coincided, in some way or other, with the phenomenon which is our defining criterion' (*Blue Book*, p. 25). For example, falling barometric pressure is a symptom of rain. Symptoms are discovered in experience; appeal to symptoms, if challenged, can be defended by appeal to empirical data; the symptoms could tomorrow cease to be reliable without the meaning of the term which refers to that of which they are currently symptoms undergoing any change, for the symptoms of the truth of a sentence are not part of the sense of the sentence. If it were raining without the barometric pressure having fallen, **[88]** we would not have doubts about the

applicability of 'It is raining.'[62] The connection between symptom and that of which it is symptomatic is inductive and contingent: the presence of symptoms is no part of the meaning of the term which refers to what the symptom points to.

At the other end of the scale there is an equally clear relation of entailment between certain sentences. 'He has never married' entails 'He is a bachelor': there is no way for a man to be a bachelor other than by having never married, or vice versa. Similarly 'It is raining' entails 'Water is falling from the sky', and 'He is in pain' entails 'He has a sensation.' This kind of definitional connection is relatively clear.

Things become less clear when attempts are made to specify a relation which falls somewhere between these two (indeed, my reference to a *scale* at the beginning of the previous paragraph was probably unfortunate). Hacker and Putnam each offers his own version of such a relation: Hacker's of course is meant as an exposition of Wittgenstein's. But it is not clear to me that the accounts which emerge of such a third relation are coherent. Let us examine them briefly.

According to Hacker the criterial relation is a novel logical relation, able to solve more than one outstanding knotty philosophical problem. It holds basically between sentences, but **[89]** we can also, as a shorthand, speak in other ways: thus we can say that a sentence is criterially related to the applicability of a word. For example, 'He screamed when a red-hot poker was laid on his

[62] It might be the case that, particularly for meteorologists, there *is* some conceptual connection between barometric pressure and rain. Not so close a connection that rain without falling pressure is unthinkable: but close enough to make falling pressure a generally expected precondition for the onset of rain, so that if the two ceased to be matched, there would be uneasiness about the concept of rain. But let us waive this possible objection here for the sake of argument: there are other cases of symptoms which involve no possible connection of this kind. To a short-sighted person, a red blob in the distance is a symptom of the approach of a bus. All buses might be painted yellow tomorrow, so that this symptomatic relation would be obliterated: but the approach of a bus would be no less the approach of a bus for that.

bare arm' is criterially related to 'He is in pain', and hence to the applicability of the word 'pain'. Even more briefly, we can say that his screaming when a red-hot poker is laid on his bare arm is a criterion for ascribing pain to him.

Before I try to characterise this criterial relation, let me set the scene by quoting Hacker's summary of his exposition of its nature: 'the criterial relation holds between sentences and derivatively between other entities. It is a fundamental semantic relation unrecognised by classical logic. It is weaker than entailment but stronger than inductive evidence. It is a relation of a priori, non-inductive, or necessarily good evidence' (Hacker, p. 293).

How is it 'stronger than inductive evidence?' Because a criterion is not disconnected from the meaning of a word in the way a symptom is. Where symptoms are discovered in experience, criteria are fixed by convention. A sentence's conventionally fixed criteria are part of the sense of the sentence, for the sense of a sentence is given by specifying the rules for its use, i.e. the criteria which justify its assertion. If *q* is a criterion for *p*, it is part of the sense of '*p*' that *q* is a priori, non-inductive, conventionally fixed evidence for the truth of '*p*'.

The cash value of this seems to be, at least in part, that while 'The barometric pressure has not fallen, but it is raining' is acceptable (though, once again, the mere fact that one wants to use 'but' and not 'and', or even to consider the two facts together as a matter of course, is significant), 'He screamed when a red-hot poker was laid on his bare arm, but he is not in pain' is not. **[90]**.

How then is the criterial relation 'weaker than entailment'? Hacker maintains that the criterial relation is evidential, which naturally means that it cannot simply be the relation of entailment. The fact that he has never married is hardly *evidence* for his being a bachelor: it *is* his being a bachelor. Having never married is the only necessary and sufficient condition for a man's being a bachelor. What then of pain, if we accept that the relation between screaming and pain is not inductive? Well, obviously screaming could not be a *necessary* condition of being in pain, for one does not always scream when one is in pain. Nor can screaming be a *sufficient*

condition of being in pain, even if it follows the laying of a red-hot poker on one's bare arm: it is still plausible to leave a loophole for not having noticed the pain at the time one screams (cf. the White Queen), or for screaming for some other reason than as an involuntary reaction to the pain, even though one does indeed feel it. So it could not be claimed that a refusal to take screaming as inductive evidence for pain commits us to saying that it is a sufficient, still less a necessary and sufficient, condition for the ascription of pain. But this is simply a difference in complexity between being a bachelor and being in pain. Someone who denies that the relation between screaming and pain is inductive must surely allow that there is a (perhaps indefinitely large) set of circumstances, the presence of an adequate group of which entails that pain is felt. What counts as an 'adequate group' of circumstances may be hard or impossible to specify, but in principle there must be some such entailment operating in our language: otherwise the 'criterial' relation collapses into one of induction. **[91]**

That this allowance must be made is denied by the criteriologist. Hacker writes (p. 291):

> Of course it is not claimed that if '*p*' is a criterion of '*q*'[63] then '*p*' is equivalent to '*q*', nor is it claimed that '*p*' entails '*q*', nor that '*q*' entails '*p*' as one of an indefinite series of disjunctive 'symptoms' for the hypothesis that *q*. Rather it is claimed that the sense of '*q*' is partially specified by the fact that the truth of '*p*' is non-inductive evidence justifying the application of '*q*'.

In reply to this sort of account, which appears frequently in Hacker, I want to say, following Wittgenstein, that repetition is not clarification.[64] Just what *is* this notion of non-inductive evidence?

[63] [Should this not read 'if *p* is a criterion of *q*'?]

[64] [Wittgenstein wrote to Bertrand Russell on 11 June 1912: 'I have just been reading a part of Moore's Principia Ethica: (now please don't be shocked) I do not like it at all. […] Moore repeats himself dozens of times, what he says in 3

'He screamed when a red-hot poker was laid on his bare arm' is not to be inductive evidence that he is in pain, nor is it to entail that he is in pain, even in conjunction with any number of other criteria: nor is 'He is in pain' to entail '*Either* he screamed when a red-hot poker was laid on his bare arm, *or* you just kicked him in the stomach, *or* he is sawing through his leg, *or* …'. Rather his screaming is a priori, non-inductive, necessarily good evidence for the truth of 'He is in pain.'

What is the cash value of this? Presumably, at least in part, that while 'He is unmarried, but he is not a bachelor' is self-contradictory, 'He screamed when a red-hot poker was laid on his bare arm, but he is not in pain' is not. But, as we have already established, nor is it acceptable to the same degree as is 'The barometric pressure has not fallen, but it is raining.' How are we to characterise the intermediate position of acceptability it supposedly occupies? It might be said that it is unacceptable in a different way from 'He is unmarried, but he is not a bachelor', in that the latter sentence is unacceptable simply in virtue of the meanings of the words used in it, while 'He screamed when a red-hot poker was laid on his bare arm, but he is not in pain' is unacceptable because to deny of someone who screams in such **[92]** circumstances that he is in pain requires an appeal to some even stronger necessarily good evidence that he is not in pain – the conjunction of the presence of the scream and the absence of pain cannot be as straightforwardly and casually observed and reported on as can the conjunction of the presence of rain and the absence of a fall in barometric pressure. And it might be said that 'He screamed when a red-hot poker was laid on his bare arm, but he is not in pain' fails to be self-contradictory in a different way from that in which 'The barometric pressure has not fallen, but it is raining' fails to be self-contradictory, in that the latter sentence fails to be self-contradictory simply in virtue of there being no semantic relation

pages could – I believe – be expressed in half a page. *Unclear* statements don't get a bit clearer by being repeated!!' *Wittgenstein in Cambridge: Letters and Documents 1911–1951*, ed. Brian McGuinness (Oxford, 2008), p. 29.

between the sentence 'The barometric pressure has fallen' and the sentence 'It is raining', whereas the former sentence fails to be self-contradictory because, although there *is* a semantic relation between 'He screamed when a red-hot poker was laid on his bare arm' and 'He is in pain', it is not so strong a semantic relation that the conjunction of the assertion of the first sentence with the denial of the second is *by itself* self-contradictory. But such accounts do not add anything to the account we already have: for they rely on the notion of necessarily good evidence, which is just what we want explained.

In short, I suspect that the notion of necessarily good evidence is empty. That which is logically necessary is not evidential, and that which is evidential is not logically necessary. A relation posited to lie between contingency and necessity requires a great deal of elucidation before it becomes a serious candidate for acceptance. What seems really to be going on here is that a strong psychological link, or an ordinary **[93]** logical link in a transitional stage of formation, is being mistaken for a non-existent 'new' type of logical link. But I shall be discussing this view in more detail in the next section, 'Teaching links and meaning'.

First, is Putnam's account of this putative intermediate relation any better? It does at least have the merit of not involving any claim about the discovery of a 'fundamental semantic relation unrecognised by classical logic': Putnam is content to use the already available notions of analyticity and syntheticity to explicate the relation with which he is concerned. I should also mention that he does not himself believe in this relation, but expounds it only as characteristic of a brand of logical behaviourism which he wishes to demolish.

Two quotes from Putnam will serve to introduce his version of the relation in question. First, on p. 3, he says that it is part of the thesis of logical behaviourism that 'there exist entailments between mind-statements and behaviour-statements; entailments that are not, perhaps, analytic in the way in which "All bachelors are unmarried" is analytic, but that nevertheless follow (in some sense)

from the meanings of mind words.' Here of course the 'in some sense' is crucial: in *what* sense?

Secondly, on p. 5, introducing his own view of the meaning of mental words in connection with 'pain', he says:

> I am inclined to say that 'pain' is a cluster-concept. That is, the application of the word 'pain' is controlled by a whole cluster of criteria, *all of which can be regarded as synthetic*. As a consequence, there is no satisfactory way of answering the question 'What does "pain" mean?' except by giving an exact synonym (e.g., 'Schmerz'); but there are a million and one different ways of saying what pain is. One can, for example, say that pain is that feeling which is normally evinced by saying 'ouch', or by wincing, or in a variety of other ways (or often not evinced at all).

[94] He explains what he means by saying that the whole cluster of criteria can be regarded as synthetic in a footnote, as follows:

> I mean not only that *each* criterion can be regarded as synthetic, but also that the cluster is *collectively* synthetic, in the sense that we are free in certain cases to say (for reason[s?] of inductive simplicity and theoretical economy) that the term applies although the whole cluster is missing. This is completely compatible with saying that the cluster serves to fix the meaning of a word. The point is that when we specify something by a cluster of indicators we assume that people will *use their brains*. That criteria may be overridden when good sense demands is the sort of thing we may regard as a 'convention associated with discourse' (Grice) rather than as something to be stipulated in connection with the individual words.

The use of the word 'fix' in this footnote is inappropriate, and the use of 'controlled' in the passage from the body of the text misleading. *If* the cluster of criteria *fixes* the meaning of the word, then we cannot correctly use the word in the absence of (some adequately representative sample of) the criteria, or deny that the word is applicable in their presence. If we *can* do either of these things, then we have, not an entailment, but a very strong inductive link.

Neither of these accounts, then, of the meaning of mental terms is clearly coherent. In reality all publicly observable phenomena are only inductive evidence for conscious mental states. Alternative accounts seem to be motivated largely by the widespread reluctance to be content with the fact that our mode of communication about subjective experience is ineluctably indirect. What is already a perfectly respectable mechanism of communication is desperately dressed up by theorists who require more convincing credentials of a word before they are prepared to grant it a use; dressed up to mimic words used of the public world. But the two kinds of word have different lives, and they cannot be assimilated. **[95]** 73

Teaching links and meaning

The upshot of this discussion is that we still have only the two logical relations of induction and entailment to play with when trying to answer the question what logical role is played by the teaching limes of sensation words. Obviously, I reject the idea that it is a simple matter of entailment. But nor do I accept that it is a simple matter of induction. I say this not because r believe in a third, 'criterial' relation such as that espoused by Hacker, but because I believe that there *is* a genuine difference between a simple inductive relation and the relation between the teaching limes for sensation words and the applicability of those words, a difference which the theory of a 'criterial' relation attempts to capture, but in the event misrepresents.

What then is this difference? Let me approach this question indirectly. The first thing to notice is that when the inessential properties of a kind of item are particularly regular, perhaps as regular as the essential properties required by the meaning of the word which describes the kind of item in question, then they tend imperceptibly to become so entrenched that their presence grows into a logical requirement too. I mention here a couple of ways in which this happens.

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

> 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.

[96] 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. The second source of the sort of meaning change under consideration is the regularity of nature. When it is a natural fact that the essential properties of a kind of thing are regularly and reliably (though inessentially) instantiated in company with other properties not definitive of the thing in question,[65] the meaning of the concept which requires the presence of the essential properties tends to change so as to incorporate as a fresh requirement that those particular inessential properties be present: the inessential thus becomes essential. Such inessential properties might be the stuff out of which common objects are made. For example, the new [1974] process of three-dimensional photography called 'holography' forms a three-dimensional image of an object in mid-air by the intersection of special laser beams. It

[65] It will often be definitive of the thing that there should be *some* such properties, but not necessarily these particular ones: a table must be made of something or other, but not necessarily of wood.

is not strictly speaking perhaps (etymologically speaking certainly) part of the meaning of 'photograph' that the image be printed on paper or projected on a screen: but our hesitation in calling a holograph a photograph may reflect the fact that our concept of a photograph has become infected by our taking for granted what is really only one sort of way, among a number of possible alternatives, of presenting a photographic image. **[97]**

Another example of such infection (though here the regularity in question is hardly the work of nature, unless it be human) might be provided by the case of the building in which an institution is housed: if Trinity College Cambridge were bombed tonight, and rebuilt in another part of Cambridge to a new design tomorrow, while the ruins of the old College were preserved as a monument after the manner of the shell of old Coventry Cathedral, we might have some hesitation about whether the identity of the College was intact. And this would not simply be because 'Trinity College' refers ambiguously to a set of buildings and a society of scholars, but because even where no such ambiguity occurs, an institution becomes associated with its architectural environment in such a way that the survival of that environment becomes to some degree definitive of the survival of the institution.

Something similar to this seems to be happening in the case of teaching links for sensation words. Since sensations are private, we have to communicate about them by means of their publicly observable causes, symptoms, manifestations, effects or whatever. These links are not always strictly part of the meaning of the sensation words, but, nature being regular, they are utterly reliable signs of the presence of the sensations with which they are correlated. So it is not surprising if an extraordinarily strong psychological connection develops between the links and the sensations, such that we feel nonplussed when we are told that all the characteristic links are present, but no sensation to match; or that a sensation is present, but no characteristic links. **[98]**

But shall we allow that this connection becomes logical? In some cases, probably it does. In some cases there is not even a question of 'becoming': the role of the concept in our lives requires

that the links be logically implicated.[66] But in most cases I don't think it is possible finally to decide, or perhaps important to do so. Meanings are in a state of flux, and at any particular moment it is hard to decide what are the precise boundaries of the field of application of a concept. What *is* important is to hold fast to saying that to the extent that the role of the links is logical, it is logical, and to the extent that it isn't, it isn't. *If* screaming becomes a logically sufficient condition for the ascription of pain (and surely it won't), then 'He screamed' *entails* 'He is in pain.' If it doesn't become such a condition, then *however strong* the psychological impulse to deny that it is possible to scream without being in pain, the relation between screaming and pain remains inductive, albeit strongly so. Agreed, there will be many difficult cases where we aren't sure whether to say that a link has encroached on a term's meaning or not: but this unclarity should be represented as such, and not disguised as a new brand of clarity by being christened 'criterial relation'.

Let us take one or two particular examples. Hare discusses whether pain is necessarily unpleasant. Obviously pain usually is unpleasant, since its function is to warn us that evasive action is called for. So it becomes psychologically very strange to hear that someone is in pain without finding it unpleasant. **[99]** And yet to the extent that 'pain' is defined in terms of the subjective quality of a sensation, we don't want to make unpleasantness a necessary condition for a sensation's being a pain. Hare escapes from this dilemma by positing two senses of 'pain' – one in which the word refers simply to a bare sensation, and another in which it refers to the same sensation, but in addition carries an implication of the co-occurrence of dislike. This is a possible move, though perhaps it would be simpler and truer to say that the presence of pain (a bare sensation) is almost irresistible (psychologically and empirically) inductive evidence for the presence of dislike. But what is *not* on is

[66] In these cases C-crude-subtle is clearly an oversimplification. But appropriate modifications will not produce any of the other versions of C which have been discussed.

to take the criteriologist's line and say that the presence of pain is 'necessarily good evidence' or somesuch for the presence of dislike. The same might be said of the relation between screaming and pain. The quotation from Hare on p. **71** expresses particularly clearly how this relation must be at least conceivably contingent.

'Necessarily good evidence' is a seductive phrase: it incorporates both the point that something is really only inductive evidence for something else, and also the point that it is such good evidence that it is beginning to infect the meaning of the word that refers to that for which it is evidence. But this infection is a sequential process through time, and is not to be characterised (incoherently) as a novel static logical relation. The notion 'necessary condition' is comprehensible; the notion 'good evidence' is comprehensible; but the notion 'necessarily good evidence' is not. Stringing words together according to the rules of surface syntax is no guarantee that the result will make sense.

Another (particularly troublesome) case: itches. The teaching links for itches is scratching, or wanting to scratch. Could there **[100]** be an itch that we didn't want to scratch? yes – a wild itch, or an itch that occurs when we are concentrating on something else that takes priority over scratching (but not concentrating so hard that we don't notice the itch at all). But even in these cases we are aware, if we stop to think about it, that the itch could be dispelled by a good scratch. Could we have a sensation of the same quality as an itch without connecting it in our minds with scratching at all? I think perhaps we could, in an absent-minded moment: I think I am sometimes conscious of an itch for a second or two before it dawns on me that the thing to do is to scratch it. Again, there might be an internal itch (though for some reason this seems an odd notion: maybe it is one of the necessary conditions for a sensation's being an itch that it be felt on the surface of the body) and thus, since scratching would be ineffective (though not perhaps rubbing or massaging or shaking or …), no thought of scratching. But suppose that these sorts of things are not possible: no itch without a psychological connection with scratching. Does this make scratching, to speak loosely, part of the meaning of 'itch'?

I think it must do, but this does not – as I shall argue in section (iii) – mean that the subjective quality of the sensation is therefore irrelevant.

One might examine a whole range of sensation terms in this way.[67] It would turn out that the connection between the links by which the use of the terms is taught and the meaning of the terms varied in strength from case to case. This is hardly surprising, because some of our sensation talk is directed towards manipulative ends, such as curing the disorder that gives rise to pain; in this case the public phenomena on which an alleviator of pain must rely in order to get any purchase, in order to have any chance of success, **[101]** are bound to loom large in the form of life in which the term 'pain' features, and so become implicated in the meaning of 'pain'. In other cases, our primary concern will be with the quality of the sensation, and here the links will be less involved in the meaning of the terms concerned. (Perhaps there are no cases where the links play simply *no* part in the meaning: but it is enough that cases of this sort should be in principle possible.) The meanings of sensation terms are variously structured and fluid, and to various degrees incorporate or come to incorporate requirements (absolute or disjunctive) that certain phenomena which act as teaching links be present before the terms are applicable: though as far as the teaching process is concerned of course, the links are of equal importance, no matter how extensive their participation in meaning. Depending on whether the links are incorporated in this way, a sudden change in the manifestations of a sensation of a certain subjective quality either would or would not entail a change in the sense of the relevant sensation term.

Putnam's discussions of multiple sclerosis and polio show particularly clearly how, in a case other than that of a sensation term, meaning can be separate from symptom links. We use the names of these and other diseases to refer to their causes (known or unknown), and not as a shorthand for (some adequate set of) their symptoms. Accordingly, when we discover that the cause of

[67] [Most of this paragraph is repetitious.]

the symptoms sometimes occurs in their absence, or that the symptoms sometimes occur without their usual cause, we have little hesitation in saying, in the former case that we nevertheless have a case of the disease, in the latter case that we don't.

As I repeatedly emphasise, it is important to keep open the possibility of talking about bare sensations, whether or not any **[102]** of our ordinary sensation terms are custom built for this purpose. Without such a way of talking, we cannot in discussion disentangle subjective and objective phenomena. For example, we could not discuss Putnam's super-super-Spartans. Putnam maintains that it makes perfectly good sense to suppose that there might be a world in which the organisms certainly had pains, just as we do, but expressed them neither through their natural expressions (thanks to an iron self-control programmed into them at birth) nor through even the most casual verbal remarks (thanks to their adherence to a strict moral custom of pain concealment, also programmed into them at birth). In view of the possibility that our human teaching links for pain may have infected the purity of the meaning of 'pain' to some degree, this claim of Putnam's may not quite go through. But it can easily be rephrased so as to be quite unexceptionable, simply by speaking not of 'pain' *tout court* (as we could do if links and meaning were safely separate), but of 'the sensation we have when we are in pain'. This periphrasis clears any semantic connection with teaching links out of the way, and enables us to discuss bare sensations quite straightforwardly. An analogous periphrasis can be used in the case of any sensational term whose meaning has got similarly entangled with its teaching links. With this small amendment, Putnam's view seems to me to be right.

Kenny's objection to Putnam's science fiction people appears to lack force. He says that 'the only reason for saying that the super-Spartans feel pain is that they say of themselves "I feel pain." Clearly, there must be some reason to think that they mean by "pain" what we mean by it; within the fiction the only such reason can be that they apply the word correctly to normal human beings in pain.' (Here, as can be seen, he is considering only super-

Spartans, who do in fact give verbal reports of their pains.) Kenny here fails **[103]** to deal with one of the central points of Putnam's article, viz. that there is another way to discover that these unusual beings feel pain, namely by examining their brains. Admittedly, as I have said, this cannot be an approved method until the correlations on which it is based have been tested by independent means: but once such tests have been successfully carried out, this technique enables us to deal with cases, just like this particular science-fictional case, where we have subjective experiences with no manifestations. Thus we may learn that super-Spartans feel pain as we do without any evidence of their being able to apply the word correctly in any circumstances. Indeed, it may well be that these beings *cannot* apply the word, or indeed any word.[68] But they may still have the same sensation as we do when we are in pain. As Putnam says, it is a mistake to confuse preconditions for talking about pain as we talk about pain (i.e. the presence of the normal teaching links) with preconditions for the existence of pain (or, in the light of what was said in the previous paragraph, for the existence of the sensation we have when we are in pain).

(iii) *What is the role of the private referent?*

It is the article of faith 'Beetles are similar' that makes the game worth playing.

Hardy (1971), p. 4

It remains to assess, in the light of what has been said in (ii), the role of the private referent in the meaning of a sensation term. First of all, itches again. Supposing we concede, as we **[104]** virtually did in the previous section, that an itch is by definition something one either wants to scratch, or something which (one realises) could be alleviated by scratching. This connection with scratching will then be a necessary condition for a sensation's being an itch. But, as I said before, it does not follow from this that the subjective quality of the sensation is irrelevant. It may *also* be a necessary condition

[68] [i.e. when they are in pain.]

for a sensation's being an itch that it feels itchy. The two conditions will be jointly sufficient,[69] but not individually so, for a sensation's being an itch. We might well want to scratch ourselves for other reasons than that we were itching: and we can imagine feeling itchy sensations independently of being aware that a scratch would be effective in removing them.

In a similar way, we can say that no matter how close a connection there is in the case of any sensation term between its meaning and the links by which its use is taught, still the mere fact that it is a *sensation* term guarantees that *one* of the necessary conditions for its correct use is that the appropriate subjectively differentiated type of private experience occurs. Agreed, a large number of our sensation terms have, because of the purposes of the discourse in which we use them, many necessary connections with public phenomena: but this does not count against the importance of our belief that the right subjective referent is there too. Sometimes this belief may play a minimal role, so that a change in the subjective referent would not be thought to unseat the meaning unduly. But we should not underestimate the general **[105]** importance of subjective regularity. It is precisely the possibility of taking it for granted that allows us to neglect the effects of its breaking down.

It is notable that in a criteriological account of the meaning of a sensation term the presence of the right subjective referent is not allowed to count as a criterion for the correct use of a sensation term. Maybe 'criterion' is so defined in its technical sense that any criterion must be a public event: but in the non-technical sense of 'criterion', it is certainly the case that in self-ascriptions of sensations the subjective quality of the experience is the criterion almost exclusively and invariably used for determining what sensation one is having. It is a pity if this ordinary-sense criterion cannot be included among the technical-sense criteria, for to refuse so to include it is to commit oneself to the view (as Hacker indeed

[69] This is not C-subtle-intermediate, for I require, while that theory does not, that itchiness be of a definite, interpersonally similar, subjective quality.

does) that self-ascriptions of experience are criterionless, and so not (among other things) either sense-bearing statements or possible candidates for being the foundations of empirical knowledge. This seems to me to be a clear mistake.

The sentence 'It looks red to me', according to Hacker (p. 308),

> presupposes that 'red' has a sense, and cannot therefore be, via the criterial nexus, part of the sense of 'It is red.' Thus although justification may in a sense regress as far as the assertion of sentences about subjective appearances, we do not find here a firm pillar upon which the superstructure of both knowledge and language may securely rest.

Now of course I can't in English understand 'It looks red to me' until I understand the meaning of 'red'. A chronology of my learning of English would relate my learning of 'It is red' before my learning of 'It looks red.' But this just happens to be the way our language is structured, and does not affect the fact that my experience of red ('sensation' of red) is a [practically] necessary **[106]** foundation[70] for my knowing how to use the word 'red'. We can imagine a language in which what we call 'looking red' is logically as well as psychologically prior to what we call 'being red'. If we learnt to speak of colours in terms of the experiences we had in their presence rather than in terms of the properties of coloured objects, we might well speak of our colour experiences, just as we

[70] Confusion of logical with temporal priority is a common mistake in philosophy. And this kind of confusion is itself only one specific example of a general tendency to deduce empirical conclusions from logical observations. Another example of the mentioned species would be to take the argument from analogy as a hypothesis about the psychology of language learning: if this were so, we would use P-predicates of ourselves in a fully-fledged way before we thought of applying them to others (cf. fig. 1, p. **9**) – and this Strawson shows to be logically impossible. In reality the argument from analogy is a retrospective justification of an assumption implicit in our actual use of P-predicates. An example of a different species would be to take Strawson's exposé of the logical structure of P-concepts as a justification of our beliefs about other minds. That a concept has a certain structure proves nothing about whether it has a genuine use: it only reveals our prejudices, which may or may not be justified.

can now speak of our pains, without reference to their origin (of which, indeed, we are often ignorant). The division between things we talk about in terms of the properties we suppose them to have independently of ourselves and things we talk about in terms of the way they impinge on our consciousness is, though certainly not arbitrary,[71] equally certainly not logically prescribed: we can imagine the dividing line being crossed in either direction. We might learn about redness via a locution like 'It φs me' (corresponding to our 'It looks red to me') and only later learn to use 'It ψs', on the grounds that it φs me (where 'It ψs' corresponds to 'It is red'). This is hard to imagine, and would probably be quite impractical, given the way the world is. But it is important that it is a logical possibility. **[107]**

[71] [How not arbitrary? Isn't the inner/outer distinction in some way *given*?]

5 SOME CONSEQUENCES OF C

If the version of C I have been sponsoring is accepted, we have a way of talking about subjective experiences as distinct from their connections with the publicly observable world: talking about them as phenomena with independently discriminable properties. This is not to deny that the language in which we so talk is logically posterior to, and so dependent upon, the language in which we talk of external matters. (Nor is it to assert this.) But this dependency, if it exists, does not affect the possibility in question.

This possibility is an extremely important and interesting one, and many further possibilities flow from it, and are dependent upon it. Opportunity is provided for speculation and further research which would on another theory all be meaningless or at least misconceived. In this final chapter I want to mention briefly one or two of the avenues opened up by an acceptance of C.[72] My remarks will be eclectic and largely programmatic.

Identity theory

One particularly fruitful possibility which remains open to those who accept C is to adopt the identity theory of sensations. Wittgenstein believed, according to Pears, that there is no way to make the description of the sensation which underlies certain teaching links more specific than '*x*' other than by reference to our attitudes to people who have *x*. But if we forge a link **[108]** between

[72] For brevity's sake I often talk in this chapter of 'C', meaning by this some suitably refined version of C-crude-subtle. This brachylogy is excusable because (*a*) versions of C-subtle are not really versions of C at all; (*b*) in the respects that matter in the present context, the differences between C-crude-crude and C-crude-subtle are not vital.

occurrences of sensations and states and/or processes of the brain, we can start to make some progress here. Obviously we will need the ordinary way of reporting and discovering sensations to get started: a brain state can in the first instance be identified as and with *x* only when the presence of *x* is established on ordinary grounds (and in an *individual* case this is of course only a necessary, not a sufficient, condition for identification). But once a sufficiently rich network of correlations has been built up, we can with confidence launch out and talk directly in neurophysiological terms, bypassing the ordinary teaching links.[73] In this way we would be able to describe sensations far more precisely than present methods 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 read off what is happening in the brain, and 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 sensations, they could surely learn to be able 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 **[109]** people's use of sensation terms suffice only to train 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 learn to use only the powers of discrimination it

[73] I know that formidable difficulties beset the coherent specification of the identity theory, let alone its acceptance, but in this programmatic context I ignore them, pointing to the benefits that would accrue if they were overcome.

needs, not all those powers which it is physically possible 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. Since we require practice to develop the skills which we are born with the capacity to manifest, untapped potential does not feature in our behavioural repertoire. After a certain age, many possibilities which were once open to us become closed: available brain space is used up in other pursuits, and cerebral pathways rigidify or atrophy. The failure to develop discriminatory potential is sometimes due to our 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 one sort of snow from all the others. It all depends on contingencies of reinforcement: where there is a need for a certain distinction to be made, so that the result of making it rightly will be significantly different from the result of making it wrongly, then, if the system at our disposal permits, we will learn to make the distinction. Differences which begin by being hardly noticeable **[110]** 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 through reinforcement is imperfectly understood, but it is a plausible enough phenomenon, and part of everyone's experience – first impressions blur distinctions; impressions of races, places, people or whatever.

In the case of the discrimination of our own sensations, the reason why we are clumsy is not that better powers of discrimination would be of no use to us – for they would; rather it is that 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 authoritatively for a large number of trials which decisions are right and which wrong, then we shall gradually grow more reliable without supervision. But if we can never be sure how we are doing, we shall 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 learn to tell reliably apart. What I am suggesting is that if, while we were trying to improve our powers of discrimination in this regard, some kindly cerebral neurophysiologist was telling us whether or not we were correctly matching sensation descriptions with brain states, we would 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.

Clearly, once the identity theory is accepted, it is no **[111]** longer the case that we have to rely on teaching links to have access to other people Is sensations, or on our own impressions to identify our own sensations. But it was precisely the necessity for the reliance on teaching links that supported the attack on C which I was trying to ward off in chapter 1. This appears to create the possibility of defending C by cutting the Gordian knot: the introduction of the identity theory would invalidate the basis of the attack on C. I did not use this manoeuvre in chapter 1 because, as I said there (p. **12**, note), to do so would beg the question as to both the possibility and the truth of C. For the identity theory depends for its coherence, let alone for its acceptance, on a prior acceptance of C. So C must first be established on other, independent, grounds. The dependence of the identity theory upon C is no argument in C's favour. Though once C is safe, and it is legitimate to enjoy the fruits of the identity theory, we may feel reassured, in a hindsightful way, in our initial partiality for C.

Intersubjective differences

Of all the speculations and questions that make sense as a consequence of adopting C – even if there is no way available for settling them, at any rate before the identity theory is adopted: and as I have said, the further we can get before adopting this theory the better – perhaps the most important are doubts about whether other people's sensations are comparable to our own either in quality or intensity (either relatively to one another or in absolute terms).[74] Before I sketch some of my reasons for thinking that doubts of this sort, to which I have alluded constantly in this thesis, have a special need to be meaningfully discussable, I will say a little more in their defence. **[112]**

According to Wisdom, this sort of doubt has no cash value. It doesn't lead us to expect anything different. Accordingly it isn't a genuine doubt. This view I reject. It is another instance of that excessive brand of verificationism against which I have already inveighed above, and which bedevils this whole area. It is motivated partly by a desire to avoid scepticism. But this can be done without such extreme measures: there is no call to confuse signs with what they signify; truth conditions with methods of verification. The common-sense model of sensation language (C) should not be carelessly abandoned in the face of a dogmatic belief that a certain unreasonably high standard of non-senseless language use must be met. Loosen the verificationist demands and retain the true model.

My objection to this point of Wisdom's can be filled out in two ways.

First: that we *expect* nothing different is not to say that we don't believe that something might *be* different. It is just that the difference we are speculating about doesn't show up, and will never show up, in observable differences. (Though, once again, on the identity theory it will. And in any case, even if we are not outright

[74] [Clarify.]

identity theorists, surely we believe that the involvement of the brain in mental states makes it exceedingly unlikely that someone might have just one 'wild' sensation which failed to match that of other people even though all the teaching links were in order. One deviance is likely to cause another, and so on until something *does* show up in the observable world. For example, there is some evidence that the danger response to red is innate: if a man saw green where we see red because his cones were misconnected, we would perhaps have a clue to this difference, even pre-physiologically, in his danger responses (a slightly **[113]** fanciful case, but it makes its point). The blithe assumption in many discussions that, for example, one's sensations of red and green, or (worse) pleasurable and painful sensations, might be straightforwardly reversed as compared with those of another person without there being any overt repercussions, let alone any effect on the distribution of other related sensations, seems extraordinary. Beliefs are not to be reduced to expectations, *in all cases*: though obviously the *general* conceptual connection between the two is extremely close.

Second: Wisdom adduces supposed parallels to the case of doubt about other people's sensations, where the doubt *can* be successfully shown to be absurd – but these cases are not really parallel. For example, someone who doubts that two nations are at war when all ordinary warlike acts are occurring has misunderstood the meaning of 'at war' (but perhaps even here there is room for doubt: perhaps the two military forces are in the pay of some wealthy film director, whose salaries are so generous that his employees are willing to risk their lives *as if* they were at war …). Similarly, perhaps it would be senseless to doubt whether someone was really dead some time after he had stopped breathing (again not a very good example: there'd be many objectors to the view that such a doubt is absurd, with many arguments at their disposal). It would be ridiculous, to take a clearer case, to doubt whether a footballer in a game of football who, without breaking any rules, had kicked the ball between the opponents' goalposts had really scored a goal or not. That is what scoring a goal *is*. In cases like

this, what is being doubted *is* perhaps reducible to a series of things all of which the doubter is accepting (though not under a description where they are explicitly constitutive of what he purports to doubt). In which case his doubt is indeed vacuous. **[114]** But doubts about the sensations of others are not so reducible: the possibility remains that his sensation is different from mine. That there is no way of checking this possibility entails that the possibility doesn't make sense only if one accepts the view that in every individual case, as well as in general, procedures for settling meaningful doubts should be at least in principle available. The intuitive view is that this is not necessary in the case of sensations.

It may seem[75] that the issues at stake in this sort of case are purely philosophical, of no importance for ordinary beliefs. But this is not so. The impression that all that is happening is that points are being made about the proper use of words like 'doubt' and 'know' is deceptive. In effect, ordinary beliefs are being questioned. The claims made are psychological: the trustworthiness of human memory is impugned, the creativity of the language learner belittled. A proposal for a revision of usage rarely amounts merely to that: for the patterns of usage reflect the patterns of belief.

I have already argued against the strong verificationist view which disallows questions of the kind being considered in this section. Compare also arguments presented by Putnam, who urges us to distinguish, in connection with related verificationist claims, the reasonable requirement that the conclusions of inductive inferences should be confirmable by other independent inferences from the unreasonable requirement that they should be non-inductively confirmable by applying a logically sufficient criterion.

To return to the point: why is the possibility of discussing intersubjective differences so important? Because such differences may have important practical consequences of a kind not normally referred to them: so that if the practical consequences are to be

[75] [How?]

[115] dealt with, the intersubjective differences must be acknowledged. Both in ordinary human social interaction and in the human sciences – psychology, sociology, anthropology and the rest – subjective experiences are given short shrift. In the former case this is because of lack of imagination, in the latter because of methodological error.

Take the former case first. It would usually be absurd to claim that an improvement in social interaction could flow from the acceptance of a philosophical thesis: particularly in the present case, since the interactors already accept the thesis at an intuitive level. Nevertheless there is a point to be made: that a particular sort of improvement depends on the truth of the thesis. So there is a practical justification for protecting the thesis against attrition by an intellectual fashion: the negative claim that social harm could issue from the rejection of a philosophical thesis is not quite so absurd.

Let me be more specific. When individuals try, singly or in groups, to understand other individuals or groups (and they don't much, in either case), it is not adequate to go purely on externals. Misunderstanding and intolerance can be based on a failure to realise that the same external state of affairs may be differently subjectively represented in different people. The facts are filtered through unique innate perceptual structures, through an individual's earlier personal history, his environmental influences, his cultural assumptions. I am not here talking simply and tritely of 'seeing the other person's point of view'. Points of view can be readily distinguished to a large degree at the objective level, by seeing what people say and do. Which arguments are dwelt on by which parties to a disagreement? Which considerations are given most weight, what general principles appealed to, what **[116]** kinds of justification offered? I am referring to something less obvious than this. The conscious experiences of two people, even in response to just the same circumstances, even when these circumstances elicit just the same behaviour, verbal and otherwise, may differ importantly, bodily sensations included. The differences may well have differential effects on future behaviour.

Because this is so, imaginative identification with others in terms of their subjective experiences can improve understanding. Necessarily, since it is so hard to discover just *how* someone else feels, this identification will take the form of allowing for the possibility of difference in experience, rather than of responding to some specific, securely identified, variation. But even this is better than nothing. Of course it is possible to know a great deal about how other people react, and about how to manipulate their reactions, without giving a thought to the quality of their experience. But if the quality of their experience is one of the sources of their reactions,[76] taking it into account will not only allow greater efficiency of manipulation, but, by bringing fuller understanding of the reasons for behaviour, will temper the antagonism and prejudice which are very near the surface even in supremely effective diplomatic relations, so long as they remain manipulative. A purely externally-minded manipulator, a social hygienist, knows that the people he is dealing with differ in their reactions in such and such ways, but secretly holds it against them, regards this as a shortcoming. They would be better off being 'normal', 'rational' – in short, like the manipulator. But if other people are regarded less as objects, their differences may be understood to proceed from different but possibly equally valid inner representations (different in quality as well as in **[117]** structure, or in spite of similarity of structure) of the environment.[77] This is surely an important difference in attitude. Not that to understand all is to forgive all: but to condemn is sometimes to fail to understand.

[76] [Is the identity theory here assumed to be correct?]

[77] Cf. Laing, p. 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.' Also p. 17: 'the relation between experience and behaviour is the stone that the builders will reject at their peril. Without it the whole structure of our theory and practice will collapse.' Strong stuff, but on the right lines.

A related failure of understanding occurs in the human sciences. These sciences reject subjective experiences as non-variables: it is not up to science to take them into account (often it is not even allowed that they *exist*). Psychology in its behaviourist forms is prominent in adopting this kind of methodological dogma. It may well be that at present the practical difficulties in the path of being sufficiently precise about subjective experiences are so great that these experiences cannot be *usefully* considered as variables. But it is an error to deny on *those* grounds that they are variables at all: for they just *are*, and important things vary as a function of them. It would be better if the disciplines in question admitted to a gap in their data, and therefore to an insecurity in their theories, rather than trying as they do to obscure the gap by phoney **[118]** methodological arguments that it is not really a gap at all.[78] That is to close off the possibility of plugging the gap when new scientific advances, especially in cerebral neurophysiology, provide the wherewithal to do so. *Already* physiology is telling psychologists a great deal, not only about intracerebral mechanics, but about the cerebral representation of subjective experiences: nevertheless some doctrinaire behaviourists still obstinately persist in drawing the boundaries of their black boxes at the interface between organism and environment. This is an arbitrary boundary, appropriate only to an earlier stage in the history of science. It is a mistake to elevate temporary practical limitations to the status of timeless tenets of good methodological doctrine.

Non-bodily experiences

In preceding chapters I have talked largely about bodily sensations. But acceptance of C allows us to consider other subjective experiences too.

[78] Cf. Poole, p. 84: 'Rather than give this world of sense-impressions the benefit of the doubt (as material for a future science, even if not the privileged geometric one) Descartes consigns it, as Galileo had done, to the refuse-heap, and there it has lain ever since.'

The concepts we use to talk about non-bodily experiences are usually such that subjective quality is not emphasised. For example, visual and auditory experiences are not normally discussed directly; more often we talk in terms of their supposed causal origins than in terms of their intrinsic subjective qualities. Indeed we may be hard put to it to recognise that any such subjective quality exists. This would be for two reasons.

First, we use visual and auditory information almost exclusively **[119]** for the purpose of interpreting our environment. This no doubt is in some way responsible for one of the most perplexing of all the phenomena of experience, the fact that we experience our environment *as external*. We think of things – colours and sounds among them – as *out there*, not in ourselves. But it would be logically possible (however difficult psychologically) to think of the world as something quite hidden from view, which we could know about to any degree only by the mediation of the experiences it caused in us. Just such a view of our situation is proposed by various 'veil of perception' doctrines which have been formulated.[79]

The second reason is that visual and auditory experiences are subjectively quite unlike bodily sensations. They are not felt as affections of any part of our bodies. (This is mysterious, since they are caused by such affections.) As a result, their subjective quality is of less interest to us than the subjective quality of sensations which do implicate our bodies. This is an oversimplification, no doubt, as anybody conscious of visual or musical beauty and ugliness will be quick to point out. But the fact remains that there is a psychical neutrality about sights and sounds that there is not about other sensations.

There are other non-bodily experiences too: the experiences which accompany or are partially constitutive of a. wide variety of mental states. They may perhaps he classed together as experiences of thought. There are 'mental' states which entail no experiences –

[79] I do not discuss here the many arguments, especially linguistic ones, against the possibility of such doctrines. But I believe they can be answered.

those for example which are particularly amenable to dispositional **[120]** analysis. And many of the experiences which accompany or partially constitute mental states are bodily: this is true particularly of emotional states, where bodily sensations are heavily implicated. But in both these cases thoughts are often involved too. And thinking is a subjective experience.

Given C, we may examine a whole range of concepts, mental and otherwise, which implicate subjective experience in one way or another, with a view to identifying the contribution of subjective experience, and the nature of the experience concerned, and assessing the importance of both. There will always be some kind of a mixture of subjective and objective criteria for the applicability of concepts of this kind: but the proportions will vary enormously, and the reasons for the variation will be instructive.

Foundations of knowledge

As mentioned at the end of chapter 4, C gives us a language in which we can discuss the claim of subjective experiences to be epistemologically basic. This claim may not be justified, but at least it can now be coherently lodged.

There is a sense in which what is claimed is bound to be true. Not that statements which report experiences are logically prior to those which report on the environment: it seems they aren't. Not that beings with behaviour and capabilities like ours are bound to be conscious: it is conceivable that this needn't be so. But given that the brain states which mediate between stimulus and response are represented in some way (not to beg any questions) in consciousness, subjective experiences are bound to be the building blocks of knowledge. They are the medium in which all knowledge is represented. Without sensory experiences there would be no occasion to talk material object language, let alone sensation **[121]** language, for there would be no access to any subjects of talk. In this sense, sense data *are* basic. Without them not. Maybe what we are aware of is always or usually heavily theory-laden: we cannot, that is, break down our experiential *Gestalten* into uninterpreted

segments, modality by modality, quale by quale. There is always some whole whose organising presence infects the quality of all its constituent parts (on this see further the immediately following section). Despite this, the fact remains that, whatever the structure of our awareness may be, we are after all *having a subjective experience*. However much knowledge may admit of extensional description, its foundations are in one way subjective. Without subjective experiences the only knowledge there could be is that parasitic knowledge we ascribe uneasily to computers.

Anti-reductionism

C might license a certain kind of attack on some forms of reductionism. Attempts to reduce various kinds of item to conglomerates of their supposedly exhaustive constituents – physical objects to sense data, minds to (dispositions to) behaviour – have been attacked on logical grounds. Proposed logical equivalences, for example, between statements about physical objects and statements about sense data are denied to hold. Or it is pointed out that physical object language is logically prior to phenomenal language. But there is also a possible *psychological* approach available, which cannot be stated without a language in which to discuss subjective experiences directly: namely to draw attention to the fact that we experience objects *as such*, an experience distinct from the sum of the experiences of the individual parts of the conglomerate to which the object is supposed to be logically reducible. **[122]** This is the kind of experience with which gestalt psychology is concerned, and it makes certain reductionist stances very implausible. Take this remark of Kovesi's (p. 19): 'We do not perceive something called "table" over and above the […] elements that have to be present in order that something should be a table.' Well, in a way we do, when we see something *as a table* (though one can see what Kovesi is getting at): the whole is experienced as other than the sum of its parts.

Neurophysiology helps to show why. Extra cells of a different order are brought into action when a number of 'parts' are presented together – cells which are never active in response to any of the parts presented individually, or even to a subgroup of such parts. In logic, a collection of properties which satisfy the necessary conceptual conditions for the presence of an object of a particular kind do not in the act of satisfaction mysteriously procreate an extra property (e.g. tableness). But in psychology a pattern of sensory stimulation adequate to generate perception of an object precisely does add to, or at least change, the subjective experience that there would be if the experiences of the uninterpreted constituent parts were simply totalled. Indeed this sort of untransformed experience does occasionally occur when it is not realised what it is that is being perceived. Any trick photograph or ambiguous figure will bear out this point.

It is a quite general feature of subjective experiences that their quality depends on the interpretation bestowed on them by the subject. We have all had experiences whose source we first misidentify, then correctly identify; and as a result we feel the experience change. This is true of bodily sensations as well as of visual experiences: the same stimulation feels quite **[123]** different depending on whether it is attributed to a clammy hand or a garment wetted by unwittingly spilt water. It may be that there is no such thing as a 'neat' sensation. Even uninterpreted sensations may be coloured by the fact that they are uninterpreted.

Consciousness

The last consequence of accepting C that I want to mention is that the problem of explaining consciousness survives in an intelligible form as a subject for deliberation.

This problem has two aspects, a scientific one and a philosophical one. At the scientific level we may ask why it is that certain arrangements of matter, and not others, should generate, or be identical with, or be accompanied by (it depends on your theory of mind) subjective experiences. Why are stones not conscious?

Why do humans need to be conscious? If sensations and other conscious states are reduced to their objective criteria these questions become far less puzzling, for they become questions about the intricacy of behavioural repertoires rather than questions about a non-behavioural phenomenon. Indeed at this level the question about stones becomes trivially equivalent to 'Why do stones not do things that stones don't do?' The issue of consciousness, properly understood, dissolves. But it is not a good idea to dissolve an issue by defining its subject matter out of existence.

The philosophical issue is closely connected with the scientific one. Apart from the relatively trivial question (though not easy to answer) of just what our criteria are for ascribing consciousness to some entities and not to others, there is the more troublesome matter of understanding consciousness even when we are secure in ascribing it. I find it hard to state the problem clearly at this level. It is not just the mind–body problem, though that might be part of it. It is the problem of a sense-conferring element **[124]** in an environment which can be regarded as quite meaningless: the problem of an awareness which finds things intelligible.[80] Groping approximations of this kind are quite inadequate, and liable to sound merely muddled and pretentiously romantic.[81] So at this stage I must be content with a manifesto. I believe that an

[80] [Answer: Finding things intelligible is just *x*, *y* … No! We don't see it that way. If we believed that we'd crack up. (The reflexivity of the conclusions of human studies: conclusions, by feedback, ruin the data on which they are based.)]

[81] In the right spirit at least, is this prayer from Simpson (1958), Act 1, Scene 2: 'PRAYER Let us throw back our heads and laugh at […] knowledge: which is an illusion caused by certain biochemical changes in the human brain structure during the course of human evolution, which had it followed another course would have produced other biochemical changes in the human brain structure, by reason of which knowledge as we now experience it would have been beyond the reach of our wildest imaginings; and by reason of which what is now beyond the reach of our wildest imaginings would have been familiar and commonplace. Let us laugh at these things. Let us laugh at thought.

'RESPONSE Which is a phenomenon like any other.'

understanding of consciousness is unavailable to conscious beings, that the problem is one of that handful located at the boundary of what it is in philosophy's power to express, let alone explain. (And it is one of philosophy's principal tasks to delineate this boundary.) If this is so, it is not surprising that philosophers are tempted by theories which dispel the problem by, in effect, ignoring it. The language in which the problem is raised only makes the issue seem problematic if it is interpreted according to the model of C: and this is just the model that tends to be banned by philosophers who find consciousness a non-issue.

It is more realistic to locate, accept and if possible state one's necessary limitations than to say that those who believe themselves to be facing up to them are the victims of some sophisticated philosophical disease. How can there not be limits to understanding? And where is one more likely to find one of them than in the self-analysis of consciousness, the agent of all analysis? A man cannot spy on himself. **[125]** As I warned, the remarks in this chapter have been programmatic. They have also been eclectic, again as advertised, for there are a number of other fruits of C which might have been discussed. Not least, the credentials of phenomenology are much strengthened. And one kind of opportunity is provided for escaping from the current plague of exclusively extensional philosophy of language. But the specification of the further implications of the theory, together with the clarification of implications already specified, I leave as a future exercise for the writer. **[126]**

BIBLIOGRAPHY

Austin, J. L. (1962), *Sense and Sensibilia* (Oxford: Clarendon Press)

Hacker, P. M. S. (1972), chapters 8 ('Private Linguists and Public Speakers') and 10 ('The Problem of Criteria') of *Insight and Illusion: Wittgenstein on Philosophy and the Metaphysics of Experience* (Oxford: Clarendon Press), pp. 215–50, 283–309

Hardy, Henry (1971), 'The Private Language Argument', unpublished

Hare, R. M. (1972), 'Pain and Evil', *Essays on the Moral Concepts* (London: Macmillan), 6, pp. 76–91

Kenny, A. J. P. (1971), reply to Putnam (1965): note 12 (p. 281) to 'The Verification Principle and the Private Language Argument (ii)', in O. R. Jones (ed.), *The Private Language Argument* (London: Macmillan), pp. 204–8 (notes pp. 280–2)

Kripke, Saul A. (1972), 'Naming and Necessity', in Donald Davidson and Gilbert Harman (eds), *Semantics of Natural Language* (Dordrecht: Reidel), pp. 253–355 (addenda pp. 763–9)

Kovesi, Julius (1967), 'Between Good and Yellow', chapter 1 of *Moral Notions* (London: Routledge & Kegan Paul), pp. 1–36

R. D. Laing (1967), *The Politics of Experience* and *The Bird of Paradise* (Harmondsworth etc.: Penguin)

Locke, Don (1968), 'The Private Language Argument', chapter 5 of *Myself and Others: A Study in Our Knowledge of Minds* (Oxford: Clarendon Press), pp. 72–109

Parfit, Derek (1971), 'Personal Identity', *Philosophical Review* 80 no. 1 (January), pp. 3–27

Pears, David (1971), 'Sensations', chapter 8 of *Wittgenstein* (London: Fontana/Collins), pp. 142–67

Poole, Roger (1972), *Towards Deep Subjectivity* (London: Allen Lane)

Putnam, Hilary (1962), 'Dreaming and "Depth Grammar"', in R. J. Butler (ed.), *Analytical Philosophy* [1st Series] (Oxford: Blackwell), pp. 211–35

Putnam, Hilary (1965) 'Brains and Behaviour', in R. J. Butler (ed.), *Analytical Philosophy*, 2nd Series (Oxford: Blackwell), 1, pp. 1–19

Putnam, Hilary (1973), 'Meaning and Reference', *Journal of Philosophy* 70 no. 19 (8 November), 699–711

Simpson, N. F. (1958), *A Resounding Tinkle: A Play* (London: Samuel French)

Strawson, P. F. (1959), 'Persons', chapter 3 of *Individuals: An Essay in Descriptive Metaphysics* (London: Methuen), pp. 87–116)

Wisdom, J. (1952) 'Other Minds I', in *Other Minds* (Oxford: Blackwell), pp. 1–34

Wittgenstein, L. (1953), *Philosophical Investigations*, trans. G. E. M. Anscombe (Oxford: Blackwell)