4

OTHER MINDS[1]

I FEEL that I agree with much, and especially with the more important parts, of what Mr. Wisdom has written, both in his present paper and in his beneficial series of articles on 'Other Minds' and other matters. I feel ruefully sure, also, that one must be at least one sort of fool to rush in over ground so well trodden by the angels. At best I can hope only to make a contribution to one part of the problem, where it seems that a little more industry still might be of service. I could only wish it was a more central part. In fact, however, I did find myself unable to approach the centre while still bogged down on the periphery. And Mr. Wisdom himself may perhaps be sympathetic towards a policy of splitting hairs to save starting them.

Mr. Wisdom, no doubt correctly, takes the 'Predicament' to be brought on by such questions as 'How do we know that another man is angry?' He also cites other forms of the question—'Do we (ever) know?', 'Can we know?', 'How can we know?' the thoughts, feelings, sensations, mind, &c., of another creature, and so forth. But it seems likely that each of these further questions is rather different from the first, which alone has been enough to keep me preoccupied, and to which I shall stick.

Mr. Wisdom's method is to go on to ask: *Is it like the way in which we know* that a kettle is boiling, or that there is a tea-party next door, or the weight of thistledown? But it seemed to me that perhaps, as he went on, he was not giving an altogether accurate account (perhaps only because too cursory a

[1] Reprinted from *Proceedings of the Aristotelian Society*, Supplementary Volume xx (1946), by courtesy of the editor.

one) of what we should say if asked 'How do you know?' these things. For example, in the case of the tea-party, to say we knew of it 'by analogy' would at best be a very sophisticated answer (and one to which some sophisticates might prefer the phrase 'by induction'), while in addition it seems incorrect because we don't, I think, claim to *know* by analogy, but only to *argue* by analogy. Hence I was led on to consider what sort of thing does actually happen when ordinary people are asked 'How do you know?'

Much depends, obviously, on the sort of item it is about which we are being asked 'How do you know?' and there are bound to be many kinds of case that I shall not cover at all, or not in detail. The sort of statement which seems simplest, and at the same time not, on the face of it, unlike 'He is angry', is such a statement as 'That is a goldfinch' ('The kettle is boiling') —a statement of particular, current, empirical fact. This is the sort of statement on making which we are liable to be asked 'How do you know?' and the sort that, at least sometimes, we say we don't know, but only believe. It may serve for a stalking-horse as well as another.

When we make an assertion such as 'There is a goldfinch in the garden' or 'He is angry', there is a sense in which we imply that we are sure of it or know it ('But I took it you *knew*', said reproachfully), though what we imply, in a similar sense and more strictly, is only that we *believe* it. On making such an assertion, therefore, we are directly exposed to the questions (1) 'Do you *know* there is?' 'Do you *know* he is?' and (2) '*How* do you know?' If in answer to the first question we reply 'Yes', we may then be asked the second question, and even the first question alone is commonly taken as an invitation to state not merely *whether* but also *how* we know. But on the other hand, we may well reply 'No' in answer to the first question: we may say 'No, but I think there is', 'No, but I believe he is'. For the implication that I know or am sure is not strict: we are not all (terribly or sufficiently) strictly brought up. If we do this, then we are exposed to the question,

which might also have been put to us without preliminaries, 'Why do you believe that?' (or 'What makes you think so?', 'What induces you to suppose so?', &c.).

There is a singular difference between the two forms of challenge: '*How* do you know?' and '*Why* do you believe?' We seem never to ask '*Why* do you know?' or '*How* do you believe?' And in this, as well as in other respects to be noticed later, not merely such other words as 'suppose' 'assume', &c., but also the expressions 'be sure' and 'be certain', follow the example of 'believe', not that of 'know'.

Either question, 'How do you know?' or 'Why do you believe?', may well be asked only out of respectful curiosity, from a genuine desire to learn. But again, they may both be asked as *pointed* questions, and, when they are so, a further difference comes out. 'How do you know?' suggests that perhaps you *don't* know it at all, whereas 'Why do you believe?' suggests that perhaps you *oughtn't* to believe it. There is no suggestion[1] that you *ought* not to know or that you *don't* believe it. If the answer to 'How do you know?' or to 'Why do you believe?' is considered unsatisfactory by the challenger, he proceeds rather differently in the two cases. His next riposte will be, on the one hand, something such as 'Then you *don't* know any such thing', or 'But that doesn't prove it: in that case you don't really know it at all', and on the other hand, something such as 'That's very poor evidence to go on: you oughtn't to believe it on the strength of that alone'.[2]

The 'existence' of your alleged belief is not challenged, but the 'existence' of your alleged knowledge *is* challenged. If we like to say that 'I believe', and likewise 'I am sure' and 'I am certain', are descriptions of subjective mental or cognitive states

[1] But in special senses and cases, there is—for example, if someone has announced some top secret information, we can ask, 'How do *you* know?', nastily.

[2] An interesting variant in the case of knowing would be 'You *oughtn't to say* (you've no business to say) you know it at all'. But of course this is only superficially similar to 'You oughtn't to believe it': you ought *to say* you believe it, if you do believe it, however poor the evidence.

or attitudes, or what not, then 'I know' is not that, or at least not merely that: it functions differently in talking.

'But of course', it will be said, ' "I know" is obviously more than that, more than a description of my own state. If I *know*, I *can't be wrong*. You can always show I don't know by showing I am wrong, or may be wrong, or that I didn't know by showing that I might have been wrong. *That's* the way in which knowing differs even from being as certain as can be.' This must be considered in due course, but first we should consider the types of answer that may be given in answer to the question 'How do you know?'

Suppose I have said 'There's a bittern at the bottom of the garden', and you ask 'How do you know?' my reply may take very different forms:

(*a*) I was brought up in the fens
(*b*) I heard it
(*c*) The keeper reported it
(*d*) By its booming
(*e*) From the booming noise
(*f*) Because it is booming.

We may say, roughly, that the first three are answers to the questions 'How do you come to know?', 'How are you in a position to know?', or 'How do *you* know?' understood in different ways: while the other three are answers to 'How can you tell?' understood in different ways. That is, I may take you to have been asking:

(1) How do I come to be in a position to know about bitterns?
(2) How do I come to be in a position to say there's a bittern here and now?
(3) How do (can) I tell bitterns?
(4) How do (can) I tell the thing here and now as a bittern?

The implication is that in order to know this is a bittern, I must have:

(1) been trained in an environment where I could become familiar with bitterns
(2) had a certain opportunity in the current case
(3) learned to recognize or tell bitterns
(4) succeeded in recognizing or telling this as a bittern.

(1) and (2) mean that my experiences must have been of certain kinds, that I must have had certain opportunities: (3) and (4) mean that I must have exerted a certain kind and amount of acumen.[1]

The questions raised in (1) and (3) concern our *past* experiences, our opportunities and our activities in learning to discriminate or discern, and, bound up with both, the correctness or otherwise of the linguistic usages we have acquired. Upon these earlier experiences depends how *well* we know things, just as, in different but cognate cases of 'knowing', it is upon earlier experience that it depends how *thoroughly* or how *intimately* we know: we know a person by sight or intimately, a town inside out, a proof backwards, a job in every detail, a poem word for word, a Frenchman when we see one. 'He doesn't know what love (real hunger) is' means he hasn't had enough experience to be able to recognize it and to distinguish it from other things slightly like it. According to how well I know an item, and according to the kind of item it is, I can recognize it, describe it, reproduce it, draw it, recite it, apply it, and so forth. Statements like 'I know *very well* he isn't angry' or 'You know *very well* that isn't calico', though of course about the current case, ascribe the excellence of the knowledge to past experience, as does the general expression 'You are old enough to know better'.[2]

[1] 'I know, I *know*, I've seen it a hundred times, don't keep on telling me' complains of a superabundance of opportunity: 'knowing a hawk from a handsaw' lays down a minimum of acumen in recognition or classification. 'As well as I know my own name' is said to typify something I *must* have experienced and *must* have learned to discriminate.

[2] The adverbs that can be inserted in 'How . . . do you know?' are few in number and fall into still fewer classes. There is practically no overlap with

By contrast, the questions raised in (2) and (4) concern the circumstances of the current case. Here we can ask 'How *definitely* do you know?' You may know it for certain, quite positively, officially, on his own authority, from unimpeachable sources, only indirectly, and so forth.

Some of the answers to the question 'How do you know?' are, oddly enough, described as 'reasons for knowing' or 'reasons to know', or even sometimes as 'reasons why I know', despite the fact that we do not ask 'Why do you know?' But now surely, according to the Dictionary, 'reasons' should be given in answer to the question 'Why?' just as we do in fact give reasons for believing in answer to the question 'Why do you believe?' However there is a distinction to be drawn here. 'How do you know that IG Farben worked for war?' 'I have every reason to know: I served on the investigating commission': here, giving my reasons for knowing is stating how I come to be in a position to know. In the same way we use the expressions 'I know *because* I saw him do it' or 'I know *because* I looked it up only ten minutes ago': these are similar to 'So it is: it *is* plutonium. How did you know?' 'I did quite a bit of physics at school before I took up philology', or to 'I ought to know: I was standing only a couple of yards away'. Reasons for *believing* on the other hand are normally quite a different affair (a recital of symptoms, arguments in support, and so forth), though there are cases where we do give as reasons for believing our having been in a position in which we could get good evidence: 'Why do you believe he was lying?' 'I was watching him very closely.'

Among the cases where we give our reasons for knowing things, a special and important class is formed by those where we cite authorities. If asked 'How do you know the election is today?', I am apt to reply 'I read it in *The Times*', and if asked 'How do you know the Persians were defeated at Marathon?' I am apt to reply 'Herodotus expressly states that they

those that can be inserted in 'How . . . do you believe?' (firmly, sincerely, genuinely, &c.).

were'. In these cases 'know' is correctly used: we know 'at second hand' when we can cite an authority who was in a position to know (possibly himself also only at second hand).[1] The statement of an authority makes me aware of something, enables me to know something, which I shouldn't otherwise have known. It is a source of knowledge. In many cases, we contrast such reasons for knowing with other reasons for believing the very same thing: 'Even if we didn't know it, even if he hadn't confessed, the evidence against him would be enough to hang him'.

It is evident, of course, that this sort of 'knowledge' is 'liable to be wrong', owing to the unreliability of human testimony (bias, mistake, lying, exaggeration, &c.). Nevertheless, the occurrence of a piece of human testimony radically alters the situation. We say 'We shall never know what Caesar's feelings were on the field of the battle of Philippi', because he did not pen an account of them: *if* he *had*, then to say 'We shall never know' won't do in the same way, even though we may still perhaps find reason to say 'It doesn't read very plausibly: we shall never *really* know the *truth*' and so on. Naturally, we are judicious: we don't say we know (at second hand) if there is any special reason to doubt the testimony: but there has to be *some* reason. It is fundamental in talking (as in other matters) that we are entitled to trust others, except in so far as there is some concrete reason to distrust them. Believing persons, accepting testimony, is the, or one main, point of talking. We don't play (competitive) games except in the faith that our opponent is trying to win: if he isn't, it isn't a game, but something different. So we don't talk with people

[1] Knowing at second hand, or on authority, is not the same as 'knowing indirectly', whatever precisely that difficult and perhaps artificial expression may mean. If a murderer 'confesses', then, whatever our opinion of the worth of the 'confession', we cannot say that 'we (only) know indirectly that he did it', nor can we so speak when a witness, reliable or unreliable, has stated that he saw the man do it. Consequently, it is not correct, either, to say that the murderer himself knows 'directly' that he did it, whatever precisely 'knowing directly' may mean.

(descriptively) except in the faith that they are trying to convey information.[1]

It is now time to turn to the question 'How can you tell?', i.e. to senses (2) and (4) of the question 'How do you know?' If you have asked 'How do you know it's a goldfinch?' then I may reply 'From its behaviour', 'By its markings', or, in more detail, 'By its red head', 'From its eating thistles'. That is, I indicate, or to some extent set out with some degree of precision, those features of the situation which enable me to recognize it as one to be described in the way I did describe it. Thereupon you may still object in several ways to my saying it's a goldfinch, without in the least 'disputing my facts', which is a further stage to be dealt with later. You may object:

(1) But goldfinches *don't* have red heads
(1*a*) But that's not a *goldfinch*. From your own description I can recognize it as a gold*crest*
(2) But that's not enough: plenty of other birds have red heads. What you say doesn't prove it. For all you know, it may be a woodpecker.

Objections (1) and (1*a*) claim that, in one way or another, I am evidently unable to recognize goldfinches. It may be (1*a*)—that I have not learned the right (customary, popular, official) name to apply to the creature ('Who taught you to use the word "goldfinch"?'):[2] or it may be that my powers of discernment, and consequently of classification, have never been brought sharply to bear in these matters, so that I remain

[1] Reliance on the authority of others is fundamental, too, in various special matters, for example, for corroboration and for the correctness of our own use of words, which we learn from others.

[2] Misnaming is not a trivial or laughing matter: If I misname I shall mislead others, and I shall also misunderstand information given by others to me. 'Of course I knew all about his condition perfectly, but I never realized that was *diabetes*: I thought it was cancer, and all the books agree that's incurable: if I'd only known it was diabetes, I should have thought of insulin at once'. Knowing *what a thing is* is, to an important extent, knowing what the name for it, and the right name for it, is.

confused as to how to tell the various species of small British bird. Or, of course, it may be a bit of both. In making this sort of accusation, you would perhaps tend not so much to use the expression 'You don't know' or 'You oughtn't to say you know' as, rather, 'But that *isn't* a goldfinch (*goldfinch*)', or 'Then you're wrong to call it a goldfinch'. But still, if asked, you would of course deny the statement that I do know it is a goldfinch.

It is in the case of objection (2) that you would be more inclined to say right out 'Then you don't know'. Because it doesn't prove it, it 's not enough to prove it. Several important points come out here:

(*a*) If you say 'That's not enough', then you must have in mind some more or less definite lack. 'To be a goldfinch, besides having a red head it must also have the characteristic eye-markings': or 'How do you know it isn't a woodpecker? Woodpeckers have red heads too'. If there is no definite lack, which you are at least prepared to specify on being pressed, then it 's silly (outrageous) just to go on saying 'That's not enough'.

(*b*) Enough is enough: it doesn't mean everything. Enough means enough to show that (within reason, and for present intents and purposes) it 'can't' be anything else, there is no room for an alternative, competing, description of it. It does *not* mean, for example, enough to show it isn't a *stuffed* goldfinch.

(*c*) '*From* its red head', given as an answer to 'How do you know?' requires careful consideration: in particular it differs very materially from '*Because* it has a red head', which is also sometimes given as an answer to 'How do you know?', and is commonly given as an answer to 'Why do you believe?' It is much more akin to such obviously 'vague' replies as 'From its markings' or 'From its behaviour' than at first appears. Our claim, in saying we know (i.e. that we can tell) is to *recognize*: and recognizing, at least in this sort of case, consists in seeing, or otherwise sensing, a feature or features which we are sure are similar to something noted (and usually named) before, on some earlier occasion in our experience. But, this that we

see, or otherwise sense, is not necessarily *describable in words*, still less describable in detail, and in non-committal words, and by anybody you please. Nearly everybody can recognize a surly look or the smell of tar, but few can describe them non-committally, i.e. otherwise than as 'surly' or 'of tar': many can recognize, and 'with certainty', ports of different vintages, models by different fashion houses, shades of green, motor-car makes from behind, and so forth, without being able to say '*how* they recognize them', i.e. without being able to 'be more specific about it'—they can only say they can tell 'by the taste', 'from the cut', and so on. So, when I say I can tell the bird 'from its red head', or that I know a friend 'by his nose', I imply that there is something *peculiar* about the red head or the nose, something peculiar to goldfinches or to him, by which you can (always) tell them or him. In view of the fewness and crudeness of the classificatory words in any language compared with the infinite number of features which are recognized, or which could be picked out and recognized, in our experience, it is small wonder that we often and often fall back on the phrases beginning with 'from' and 'by', and that we are not able to *say*, further and precisely, *how* we can tell. Often we know things quite well, while scarcely able at all to say 'from' what we know them, let alone what there is so very special about them. Any answer beginning 'From' or 'By' has, intentionally, this saving 'vagueness'. But on the contrary, an answer beginning 'Because' is dangerously definite. When I say I know it's a goldfinch 'Because it has a red head', that implies that all I have noted, or needed to note, about it is that its head is red (nothing special or peculiar about the shade, shape, &c. of the patch): so that I imply that there is no other small British bird that has any sort of red head except the goldfinch.

(*d*) Whenever I say I know, I am always liable to be taken to claim that, in a certain sense appropriate to the kind of statement (and to present intents and purposes), I am able to *prove* it. In the present, very common, type of case, 'proving' seems to mean stating what are the features of the current case

which are enough to constitute it one which is correctly describable in the way we have described it, and not in any other way relevantly variant. Generally speaking, cases where I can 'prove' are cases where we use the 'because' formula: cases where we 'know but can't prove' are cases where we take refuge in the 'from' or 'by' formula.

I believe that the points so far raised are those most genuinely and normally raised by the question 'How do you know?' But there are other, further, questions sometimes raised under the same rubric, and especially by philosophers, which may be thought more important. These are the worries about 'reality' and about being 'sure and certain'.

Up to now, in challenging me with the question 'How do you know?', you are not taken to have *queried my credentials as stated*, though you have asked what they were: nor have you *disputed my facts* (the facts on which I am relying to prove it is a goldfinch), though you have asked me to detail them. It is this further sort of challenge that may now be made, a challenge as to the *reliability* of our alleged 'credentials' and our alleged 'facts'. You may ask:

(1) But do you know it's a *real* goldfinch? How do you know you're not dreaming? Or after all, mightn't it be a stuffed one? And is the head really red? Couldn't it have been dyed, or isn't there perhaps an odd light reflected on it?

(2) But are you certain it's the *right* red for a goldfinch? Are you quite sure it isn't too orange? Isn't it perhaps rather too strident a note for a bittern?

These two sorts of worry are distinct, though very probably they can be combined or confused, or may run into one another: e.g. 'Are you sure it's really red?' may mean 'Are you sure it isn't orange?' or again 'Are you sure it isn't just the peculiar light?'

1. *Reality*

If you ask me, 'How do you know it's a real stick?' 'How do you know it's really bent?' ('Are you sure he's really angry?'),

then you are querying my credentials or my facts (it's often uncertain which) in a certain special way. In various *special, recognized* ways, depending essentially upon the nature of the matter which I have announced myself to know, either my current experiencing or the item currently under consideration (or uncertain which) may be abnormal, *phoney*. Either I myself may be dreaming, or in delirium, or under the influence of mescal, &c.: or else the item may be stuffed, painted, dummy, artificial, trick, freak, toy, assumed, feigned, &c.: or else again there's an uncertainty (it's left open) whether *I* am to blame or *it* is—mirages, mirror images, odd lighting effects, &c.

These doubts are all to be allayed by means of recognized procedures (more or less roughly recognized, of course), appropriate to the particular type of case. There are recognized ways of distinguishing between dreaming and waking (how otherwise should we know how to use and to contrast the words?), and of deciding whether a thing is stuffed or live, and so forth. The doubt or question 'But is it a *real* one?' has always (*must* have) a special basis, there must be some 'reason for suggesting' that it isn't real, in the sense of some specific way, or limited number of specific ways, in which it is suggested that this experience or item may be phoney. Sometimes (usually) the context makes it clear what the suggestion is: the goldfinch might be stuffed but there's no suggestion that it's a mirage, the oasis might be a mirage but there's no suggestion it might be stuffed. If the context doesn't make it clear, then I am entitled to ask 'How do you mean? Do you mean it may be stuffed or what? *What are you suggesting?*' The wile of the metaphysician consists in asking 'Is it a real table?' (a kind of object which has no obvious way of being phoney) and not specifying or limiting what may be wrong with it, so that I feel at a loss 'how to prove' it *is* a real one.[1] It is the use of the word 'real' in this

[1] Conjurers, too, trade on this. 'Will some gentleman kindly satisfy himself that this is a perfectly ordinary hat?' This leaves us baffled and uneasy: sheepishly we agree that it seems all right, while conscious that we have not the least idea what to guard against.

manner that leads us on to the supposition that 'real' has a single meaning ('the real world' 'material objects'), and that a highly profound and puzzling one. Instead, we should insist always on specifying with what 'real' is being contrasted—'not what' I shall have to show it is, in order to show it is 'real': and then usually we shall find some specific, less fatal, word, appropriate to the particular case, to substitute for 'real'.

Knowing it's a 'real' goldfinch isn't in question in the ordinary case when I say I know it's a goldfinch: reasonable precautions only are taken. But when it *is* called in question, in *special* cases, then I make sure it's a real goldfinch in ways essentially similar to those in which I made sure it was a goldfinch, though corroboration by other witnesses plays a specially important part in some cases. Once again the precautions cannot be more than reasonable, relative to current intents and purposes. And once again, in the special cases just as in the ordinary cases, two further conditions hold good:

(*a*) I don't by any means *always* know whether it's one or not. It may fly away before I have a chance of testing it, or of inspecting it thoroughly enough. This is simple enough: yet some are prone to argue that because I *sometimes* don't know or can't discover, I *never* can.

(*b*) 'Being sure it's real' is no more proof against miracles or outrages of nature than anything else is or, *sub specie humanitatis*, can be. If we have made sure it's a goldfinch, and a real goldfinch, and then in the future it does something outrageous (explodes, quotes Mrs. Woolf, or what not), we don't say we were wrong to say it was a goldfinch, *we don't know what to say*. Words literally fail us: 'What would you have said?' 'What are we to say now?' 'What would *you* say?' When I have made sure it's a real goldfinch (not stuffed, corroborated by the disinterested, &c.) then I am *not* 'predicting' in saying it's a real goldfinch, and in a very good sense I can't be proved wrong whatever happens. It seems a serious mistake to suppose that language (or most language, language about real things) is 'predictive' in such a way that the future can

always prove it wrong. What the future *can* always do, is to make us *revise our ideas* about goldfinches or real goldfinches or anything else.

Perhaps the normal procedure of language could be schematized as follows. First, it is arranged that, on experiencing a complex of features C, then we are to say 'This is C' or 'This is a C'. Then subsequently, the occurrence either of the whole of C or of a significant and characteristic part of it is, on one or many occasions, accompanied or followed in definite circumstances by another special and distinctive feature or complex of features, which makes it seem desirable to revise our ideas: so that we draw a distinction between 'This looks like a C, but in fact is only a dummy, &c.' and 'This is a real C (live, genuine, &c.)'. *Henceforward*, we can only ascertain that it's a *real* C by ascertaining that the special feature or complex of features is present in the appropriate circumstances. The old expression 'This is a C' will tend as heretofore to fail to draw any distinction between 'real, live, &c.' and 'dummy, stuffed, &c.' If the special distinctive feature is one which does not have to manifest itself in *any* definite circumstances (on application of some specific test, after some limited lapse of time, &c.) then it is not a suitable feature on which to base a distinction between 'real' and 'dummy, imaginary, &c.' All we can then do is to say 'Some Cs are and some aren't, some do and some don't: and it may be very interesting or important whether they are or aren't, whether they do or don't, but they're all Cs, real Cs, just the same'.[1] Now if the special feature is one which must appear in (more or less) definite circumstances, then 'This is a real C' is not necessarily predictive: we can, in favourable cases, make sure of it.[2]

[1] The awkwardness about some snarks being boojums.

[2] Sometimes, on the basis of the new special feature, we distinguish, not between 'Cs' and 'real Cs', but rather between Cs and Ds. There is a reason for choosing the one procedure rather than the other: all cases where we use the 'real' formula exhibit (complicated and serpentine) likenesses, as do all cases where we use 'proper', a word which behaves in many ways like 'real', and is no less nor more profound.

2. *Sureness and Certainty*

The other way of querying my credentials and proofs ('Are you sure it's the *right* red?') is quite different. Here we come up against Mr. Wisdom's views on 'the peculiarity of a man's knowledge of his own sensations', for which he refers us to 'Other Minds VII' (*Mind*, vol. lii, N.S., no. 207), a passage with which I find I disagree.

Mr. Wisdom there says that, excluding from consideration cases like 'being in love' and other cases which 'involve prediction', and considering statements like 'I am in pain' which, in the requisite sense, do *not* involve prediction, then a man *cannot* 'be wrong' in making them, in the most favoured sense of being wrong: that is, though it is of course possible for him to *lie* (so that 'I am in pain' may be false), and though it is also possible for him to *misname*, i.e. to use the word 'pawn', say, instead of 'pain', which would be liable to mislead others but would not mislead himself, either because he regularly uses 'pawn' for 'pain' or because the use was a momentary aberration, as when I call John 'Albert' while knowing him quite well to be John—though it is possible for him to be 'wrong' in these two senses, it is not possible for him to be wrong in the most favoured sense. He says again that, with this class of statement (elsewhere called 'sense-statements'), to know directly that one is in pain is 'to say that one is, and to say it on the basis of being in pain': and again, that the peculiarity of sense-statements lies in the fact that 'when they are correct and made by X, then X knows they are correct'.

This seems to me mistaken, though it is a view that, in more or less subtle forms, has been the basis of a very great deal of philosophy. It is perhaps the original sin (Berkeley's apple, the tree in the quad) by which the philosopher casts himself out from the garden of the world we live in.

Very clearly detailed, this is the view that, at least and only in a certain favoured type of case, I can 'say what I see (or otherwise sense)' almost quite literally. On this view, if I were to say 'Here is something red', then I might be held to imply or

to state that it is really a red thing, a thing which would appear red in a standard light, or to other people, or tomorrow too, and perhaps even more besides: all of which 'involves prediction' (if not also a metaphysical substratum). Even if I were to say 'Here is something which looks red', I might still be held to imply or to state that it looks red to others also, and so forth. If, however, I confine myself to stating 'Here is something that looks red to me now', then at last I can't be wrong (in the most favoured sense).

However, there is an ambiguity in 'something that looks red to me now'. Perhaps this can be brought out by italics, though it is not really so much a matter of emphasis as of tone and expression, of confidence and hesitancy. Contrast 'Here is something that (definitely) *looks to me* (anyhow) red' with 'Here is something that looks to me (something like) *red* (I should say)'. In the former case I am quite confident that, however it may look to others, whatever it may 'really be', &c., it certainly does look red to me at the moment. In the other case I'm not confident at all: it looks reddish, but I've never seen anything quite like it before, I can't quite describe it—or, I'm not very good at recognizing colours, I never feel quite happy about them, I've constantly been caught out about them. Of course, this sounds silly in the case of 'red': red is so *very* obvious, we all know red when we see it, it's *unmistakable*.[1] Cases where we should not feel happy about red are not easy (though not impossible) to find. But take 'magenta': 'It looks rather like magenta to me—but then I wouldn't be too sure about distinguishing magenta from mauve or from heliotrope. Of course I know in a way it's purplish, but I don't really know whether to say it's magenta or not: I just can't be sure.' Here, I am not interested in ruling out consideration of how it looks to others (looks *to me*) or considerations about what its *real* colour is (*looks*): what I am ruling out is *my being sure or certain* what it looks to me. Take tastes, or take sounds:

[1] And yet she always *thought* his shirt was white until she saw it against Tommy's Persil-washed one.

these are so much better as examples than colours, because we never feel so happy with our other senses as with our eyesight. Any description of a taste or sound or smell (or colour) or of a feeling, involves (is) saying that it is like one or some that we have experienced before: any descriptive word is classificatory, involves recognition and in that sense memory, and only when we use such words (or names or descriptions, which come down to the same) are we knowing anything, or believing anything. But memory and recognition are often uncertain and unreliable.

Two rather different ways of being hesitant may be distinguished.

(*a*) Let us take the case where we are tasting a certain taste. We may say 'I simply don't know what it is: I've never tasted anything remotely like it before. . . . No, it's no use: the more I think about it the more confused I get: it's perfectly distinct and perfectly distinctive, quite unique in my experience'. This illustrates the case where I can find nothing in my past experience with which to compare the current case: I'm certain it's not appreciably like anything I ever tasted before, not sufficiently like anything I know to merit the same description. This case, though distinguishable enough, shades off into the more common type of case where I'm not quite certain, or only fairly certain, or practically certain, that it's the taste of, say, laurel. In all such cases, I am endeavouring to recognize the current item by searching in my past experience for something like it, some likeness in virtue of which it deserves, more or less positively, to be described by the same descriptive word:[1] and I am meeting with varying degrees of success.

(*b*) The other case is different, though it very naturally combines itself with the first. Here, what I try to do is to *savour* the current experience, to *peer* at it, to sense it vividly. I'm not sure it *is* the taste of pineapple: isn't there perhaps just *some-*

[1] Or, of course, related to it in some other way than by 'similarity' (in any ordinary sense of 'similarity'), which is yet sufficient reason for describing it by the same word.

thing about it, a tang, a bite, a lack of bite, a cloying sensation, which isn't *quite* right for pineapple? Isn't there perhaps just a peculiar hint of green, which would rule out mauve and would hardly do for heliotrope? Or perhaps it is faintly odd: I must look more intently, scan it over and over: maybe just possibly there is a suggestion of an unnatural shimmer, so that it doesn't look quite like ordinary water. There is a lack of sharpness in what we actually sense, which is to be cured not, or not merely, by thinking, but by acuter discernment, by sensory discrimination (though it is of course true that thinking of other, and more pronounced, cases in our past experience can and does assist our powers of discrimination).[1]

Cases (*a*) and (*b*) alike, and perhaps usually together, lead to our being not quite sure or certain what it is, what to say, how to describe it: what our feelings really are, whether the tickling is painful exactly, whether I'm really what you'd call angry with him or only something rather like it. The hesitation is of course, in a sense, over misnaming: but I am not so much or merely worried about possibly misleading others as about misleading myself (the most favoured sense of being wrong). I should suggest that the two expressions 'being certain' and 'being sure', though from the nature of the case they are often used indiscriminately, have a tendency to refer to cases (*a*) and (*b*) respectively. 'Being certain' tends to indicate confidence in our memories and our past discernment, 'being sure' to indicate confidence in the current perception. Perhaps this comes out in our use of the concessives 'to be sure' and 'certainly', and in our use of such phrases as 'certainly not 'and 'surely not'. But it may be unwise to chivvy language beyond the coarser nuances.

It may be said that, even when I don't know exactly how to describe it, I nevertheless *know* that I *think* (and roughly how confidently I think) it is mauve. So I do know *something*. But

[1] This appears to cover cases of dull or careless or uninstructed perception, as opposed to cases of diseased or drugged perception.

this is irrelevant: I *don't* know it's mauve, that it definitely looks to me now mauve. Besides, there are cases where I really don't know what I think: I'm completely baffled by it.

Of course, there are any number of 'sense-statements' about which I can be, and am, completely sure. In ordinary cases ordinary men are nearly always certain when a thing looks red (or reddish, or anyhow reddish rather than greenish), or when they're in pain (except when that's rather difficult to say, as when they're being tickled): in ordinary cases an expert, a dyer or a dress designer, will be quite sure when something looks (to him in the present light) reseda green or nigger brown, though those who are not experts will not be so sure. Nearly always, if not quite always, we can be quite, or pretty, sure if we take refuge in a sufficiently *rough* description of the sensation: roughness and sureness tend to vary inversely. But the less rough descriptions, just as much as the rough, are all 'sense-statements'.

It is, I think, the problems of sureness and certainty, which philosophers tend (if I am not mistaken) to neglect, that have considerably exercised scientists, while the problem of 'reality', which philosophers have cultivated, does not exercise them. The whole apparatus of measures and standards seems designed to combat unsureness and uncertainty, and concomitantly to increase the possible precision of language, which, in science, pays. But for the words 'real' and 'unreal' the scientist tends to substitute, wisely, their cash-value substitutes, of which he invents and defines an increasing number, to cover an increasing variety of cases: he doesn't ask 'Is it real?' but rather 'Is it denatured?' or 'Is it an allotropic form?' and so on.

It is not clear to me what the class of sense-statements is, nor what its 'peculiarity' is. Some who talk of sense-statements (or sense data) appear to draw a distinction between talking about simple things like red or pain, and talking about complicated things like love or tables. But apparently Mr. Wisdom does not, because he treats 'This looks to me now like a man eating poppies' as in the same case with 'This looks to me now red'.

In this he is surely right: a man eating poppies may be more 'complex' to recognize, but it is often not appreciably more difficult than the other. But if, again, we say that non-sense-statements are those which involve 'prediction', why so? True, if I say, 'This is a (real) oasis' without first ascertaining that it's not a mirage, then I do chance my hand: but if I *have* ascertained that it's not, and can recognize for sure that it isn't (as when I am drinking its waters), then surely I'm not chancing my hand any longer. I believe, of course, that it will continue to perform as (real) oases normally do: but if there's a *lusus naturae*, a miracle, and it doesn't, that wouldn't mean I was wrong, previously, to call it a real oasis.

With regard to Mr. Wisdom's own chosen formulae, we have seen already that it can't be right to say that the peculiarity of sense-statements is that 'when they are correct, and made by X, then X knows they are correct': for X may *think*, without much confidence, that it tastes to him like Lapsang, and yet be far from certain, and then subsequently become certain, or more certain, that it did or didn't. The other two formulae were: 'To know that one is in pain is to say that one is and to say it on the basis of being in pain' and that the only mistake possible with sense-statements is typified by the case where 'knowing him to be Jack I call him "Alfred", thinking his name is Alfred or not caring a damn what his name is'. The snag in both these lies in the phrases 'on the basis of being in pain' and 'knowing him to be Jack'. 'Knowing him to be Jack' means that I have recognized him as Jack, a matter over which I may well be hesitant and/or mistaken: it is true that I needn't recognize him *by name* as 'Jack' (and hence I may call him 'Alfred'), but at least I must be recognizing him correctly as, for instance, the man I last saw in Jerusalem, or else I *shall* be misleading *myself*. Similarly, if 'on the basis of being in pain' only means 'when I am (what would be correctly described as) in pain', then something more than merely *saying* 'I'm in pain' is necessary for knowing I'm in pain: and this something more, as it involves recognition, may be hesitant and/or mistaken,

though it is of course unlikely to be so in a case so comparatively obvious as that of pain.

Possibly the tendency to overlook the problems of recognition is fostered by the tendency to use a direct object after the word *know*. Mr. Wisdom, for example, confidently uses such expressions as 'knowing the feelings of another (his mind, his sensations, his anger, his pain) in the way that *he* knows them'. But, although we do correctly use the expressions 'I know your feelings on the matter' or 'He knows his own mind' or (archaically) 'May I know your mind?', these are rather special expressions, which do not justify any general usage. 'Feelings' here has the sense it has in 'very strong feelings' in favour of or against something: perhaps it means 'views' or 'opinions' ('very decided opinions'), just as 'mind' in this usage is given by the Dictionary as equivalent to 'intention' or 'wish'. To extend the usage uncritically is somewhat as though, on the strength of the legitimate phrase 'knowing someone's tastes', we were to proceed to talk of 'knowing someone's sounds' or 'knowing someone's taste of pineapple'. If, for example, it is a case of *physical* feelings such as fatigue, we do not use the expression 'I know your feelings'.

When, therefore, Mr. Wisdom speaks generally of 'knowing his sensations', he presumably means this to be equivalent to 'knowing *what* he is seeing, smelling, &c.', just as 'knowing the winner of the Derby' means 'knowing *what won* the Derby'. But here again, the expression 'know what' seems sometimes to be taken, unconsciously and erroneously, to lend support to the practice of putting a direct object after *know*: for 'what' is liable to be understood as a relative, = 'that which'. This is a grammatical mistake: 'what' *can* of course be a relative, but in 'know what you feel' and 'know what won' it is an interrogative (Latin *quid*, not *quod*). In this respect, 'I can smell what he is smelling' differs from 'I can know what he is smelling'. 'I know what he is feeling' is not 'There is an x which both I know and he is feeling', but 'I know the answer to the question "What is he feeling?"' And similarly with 'I know what I am

feeling': this does *not* mean that there is something which I am *both knowing and feeling.*

Expressions such as 'We don't know another man's anger in the way he knows it' or 'He knows his pain in a way we can't' seem barbarous. The man doesn't 'know his pain': he feels (not knows) what he recognizes as, or what he knows to be, anger (not his anger), and he knows that he is feeling angry. Always assuming that he does recognize the feeling, which in fact, though feeling it acutely, he may not: 'Now I know what it was, it was jealousy (or gooseflesh or angina). At the time I did not know at all what it was, I had never felt anything quite like it before: but since then I've got to know it quite well'.[1]

Uncritical use of the direct object after *know* seems to be one thing that leads to the view that (or to talking as though) sensa, that is things, colours, noises, and the rest, speak or are labelled by nature, so that I can literally *say* what (that which) I *see*: it pipes up, or I read it off. It is as if sensa were *literally* to 'announce themselves' or to 'identify themselves', in the way we indicate when we say 'It presently identified itself as a particularly fine white rhinoceros'. But surely this is only a manner of speaking, a reflexive idiom in which the French, for example, indulge more freely than the English: sensa are dumb, and only previous experience enables *us* to identify them. If we choose to say that they 'identify themselves' (and certainly 'recognizing' is not a highly voluntary activity of ours), then it must be admitted that they share the birthright of all speakers, that of speaking unclearly and untruly.

If I know I can't be wrong

One final point about 'How do you know?', the challenge to the user of the expression 'I know', requires still to be brought

[1] There are, of course, legitimate uses of the direct object after *know*, and of the possessive pronoun before words for feelings. 'He knows the town well', 'He has known much suffering', 'My old vanity, how well I know it!'—even the pleonastic 'Where does he feel his (= the) pain?' and the educative tautology '*He* feels *his* pain'. But none of these really lends support to the metaphysical 'He knows his pain (in a way we can't)'.

out by consideration of the saying that 'If you know you can't be wrong'. Surely, if what has so far been said is correct, then we are often right to say we *know* even in cases where we turn out subsequently to have been mistaken—and indeed we seem always, or practically always, liable to be mistaken.

Now, we are perfectly, and should be candidly, aware of this liability, which does not, however, transpire to be so very onerous in practice. The human intellect and senses are, indeed, *inherently* fallible and delusive, but not by any means *inveterately* so. Machines are inherently liable to break down, but good machines don't (often). It is futile to embark on a 'theory of knowledge' which denies this liability: such theories constantly end up by admitting the liability after all, and denying the existence of 'knowledge'.

'When you know you can't be wrong' is perfectly good sense. You are prohibited from saying 'I know it is so, but I may be wrong', just as you are prohibited from saying 'I promise I will, but I may fail'. If you are aware you may be mistaken, you ought not to say you know, just as, if you are aware you may break your word, you have no business to promise. But of course, being aware that you may be mistaken doesn't mean merely being aware that you are a fallible human being: it means that you have some concrete reason to suppose that you may be mistaken in this case. Just as 'but I may fail' does not mean merely 'but I am a weak human being' (in which case it would be no more exciting than adding 'D.V.'): it means that there is some concrete reason for me to suppose that I shall break my word. It is naturally *always* possible ('humanly' possible) that I may be mistaken or may break my word, but that by itself is no bar against using the expressions 'I know' and 'I promise' as we do in fact use them.

At the risk (long since incurred) of being tedious, the parallel between saying 'I know' and saying 'I promise' may be elaborated.[1]

[1] It is the use of the expressions 'I know' and 'I promise' (first person singular, present indicative tense) alone that is being considered. 'If I knew,

When I say 'S is P', I imply at least that I believe it, and, if I have been strictly brought up, that I am (quite) sure of it: when I say 'I shall do A', I imply at least that I hope to do it, and, if I have been strictly brought up that I (fully) intend to. If I only believe that S is P, I can add 'But of course I may (very well) be wrong:' if I only hope to do A, I can add 'But of course I may (very well) not'. When I only believe or only hope, it is recognized that further evidence or further circumstances are liable to make me change my mind. If I say 'S is P' when I don't even believe it, I am lying: if I say it when I believe it but am not sure of it, I may be misleading but I am not exactly lying. If I say 'I shall do A' when I have not even any hope, not the slightest intention, of doing it, then I am deliberately deceiving: if I say it when I do not fully intend to, I am misleading but I am not deliberately deceiving in the same way.

But now, when I say 'I promise', a new plunge is taken: I have not merely announced my intention, but, by using this formula (performing this ritual), I have bound myself to others, and staked my reputation, in a new way. Similarly, saying 'I know' is taking a new plunge. But it is *not* saying 'I have performed a specially striking feat of cognition, superior, in the same scale as believing and being sure, even to being merely quite sure': for there *is* nothing in that scale superior to being quite sure. Just as promising is not something superior, in the same scale as hoping and intending, even to merely fully intending: for there *is* nothing in that scale superior to fully intending. When I say 'I know', I *give others my word*: I *give others my authority for saying* that 'S is P'.

I can't have been wrong' or 'If she knows she can't be wrong' are not worrying in the way that 'If I ("you") know I ("you") can't be wrong' is worrying. Or again, 'I promise' is quite different from 'he promises': if I say 'I promise', I don't say I *say* I promise, I *promise*, just as if he says he promises, he doesn't say he says he promises, he promises: whereas if I say 'he promises', I do (only) say he *says* he promises—in the other 'sense' of 'promise', the 'sense' in which *I* say *I* promise, only *he* can say he promises. I *describe* his promising, but I *do* my own promising and he must do *his* own.

When I have said only that I am sure, and prove to have been mistaken, I am not liable to be rounded on by others in the same way as when I have said 'I know'. I am sure *for my part*, you can take it or leave it: accept it if you think I'm an acute and careful person, that's your responsibility. But I don't know 'for my part', and when I say 'I know' I don't mean you can take it or leave it (though of course you *can* take it or leave it). In the same way, when I say I fully intend to, I do so for my part, and, according as you think highly or poorly of my resolution and chances, you will elect to act on it or not to act on it: but if I say I promise, you are *entitled* to act on it, whether or not you choose to do so. If I have said I know or I promise, you insult me in a special way by refusing to accept it. We all *feel* the very great difference between saying even 'I'm *absolutely* sure' and saying 'I know': it is like the difference between saying even 'I firmly and irrevocably intend' and 'I promise'. If someone has promised me to do A, then I am entitled to rely on it, and can myself make promises on the strength of it: and so, where someone has said to me 'I know', I am entitled to say *I* know too, at second hand. The right to say 'I know' is transmissible, in the sort of way that other authority is transmissible. Hence, if I say it lightly, I may be *responsible* for getting *you* into trouble.

If you say you *know* something, the most immediate challenge takes the form of asking 'Are you in a position to know?': that is, you must undertake to show, not merely that you are sure of it, but that it is within your cognisance. There is a similar form of challenge in the case of promising: fully intending is not enough—you must also undertake to show that 'you are in a position to promise', that is, that it is within your power. Over these points in the two cases parallel series of doubts are apt to infect philosophers, on the ground that I cannot foresee the future. Some begin to hold that I should never, or practically never, say I know anything—perhaps only what I am sensing at this moment: others, that I should never, or practically never, say I promise—perhaps only what is actually within

my power at this moment. In both cases there is an obsession: if I know *I can't be wrong*, so I can't have the right to say I know, and if I promise *I can't fail*, so I can't have the right to say I promise. And in both cases this obsession fastens on my inability to make *predictions* as the root of the matter, meaning by predictions claims to know the future. But this is doubly mistaken in both cases. As has been seen, we may be perfectly justified in saying we know or we promise, in spite of the fact that things 'may' turn out badly, and it's a more or less serious matter for us if they do. And further, it is overlooked that the conditions which must be satisfied if I am to show that a thing is within my cognisance or within my power are conditions, not about the future, but about *the present and the past*: it is not demanded that I do more than *believe* about the future.[1]

We feel, however, an objection to saying that 'I know' performs the same sort of function as 'I promise'. It is this. Supposing that things turn out badly, then we say, on the one hand 'You're proved wrong, so you *didn't* know', but on the other hand 'You've failed to perform, although you *did* promise'. I believe that this contrast is more apparent than real. The sense in which you 'did promise' is that you did *say* you promised (did say 'I promise'): and you did *say* you knew. That is the gravamen of the charge against you when you let us down, after we have taken your word. But it may well transpire that you never fully intended to do it, or that you had concrete reason to suppose that you wouldn't be able to do it (it might even be manifestly impossible), and in another 'sense' of promise you *can't* then have promised to do it, so that you *didn't* promise.

Consider the use of other phrases analogous to 'I know' and 'I promise'. Suppose, instead of 'I know', I had said 'I swear': in that case, upon the opposite appearing, we should say, exactly as in the promising case, 'You *did* swear, but you were

[1] If 'Figs never grow on thistles' is taken to mean 'None ever have and none ever will', then it is implied that I *know* that none ever have, but only that I *believe* that none ever will.

wrong'. Suppose again that, instead of 'I promise', I had said 'I guarantee' (e.g. to protect you from attack): in that case, upon my letting you down, you can say, exactly as in the knowing case 'You *said* you guaranteed it, but you *didn't* guarantee it'.[1] Can the situation perhaps be summed up as follows? In these 'ritual' cases, the approved case is one where *in the appropriate circumstances*, I say a certain formula: e.g. 'I do' when standing, unmarried or a widower, beside woman, unmarried or a widow and not within the prohibited degrees of relationship, before a clergyman, registrar, &c., or 'I give' when it is mine to give, &c., or 'I order' when I have the authority to, &c. But now, if the situation transpires to have been in some way not orthodox (I was already married: it wasn't mine to give: I had no authority to order), then we tend to be rather hesitant about how to put it, as heaven was when the saint blessed the penguins. We call the man a bigamist, but his second marriage was not a marriage, is null and void (a useful formula in many cases for avoiding saying either 'he did' or 'he didn't'): he did 'order' me to do it, but, having no authority over me, he *couldn't* 'order' me: he did warn me it was going to charge, but it wasn't or anyway I knew much more about it than he did, so in a way he couldn't warn me, didn't warn me.[2] We hesitate between 'He didn't order me', 'He had no right to order me', 'He oughtn't to have said he ordered me', just as we do between 'You didn't know', 'You can't have known', 'You had no right to say you knew' (these perhaps having slightly different nuances, according to what precisely it is that has gone wrong). But the essential factors

[1] 'Swear', 'guarantee', 'give my word', 'promise', all these and similar words cover cases both of 'knowing' and of 'promising', thus suggesting the two are analogous. Of course they differ subtly from each other; for example, *know* and *promise* are in a certain sense 'unlimited' expressions, while when I swear I swear *upon* something, and when I guarantee I guarantee that, upon some adverse and more or less to be expected circumstance arising, I will take *some more or less definite action* to nullify it.

[2] 'You can't warn someone of something that isn't going to happen' parallels 'You can't know what isn't true'.

are (*a*) You said you knew: you said you promised (*b*) You were mistaken: you didn't perform. The hesitancy concerns only the precise way in which we are to round on the original 'I know' or 'I promise'.

To suppose that 'I know' is a descriptive phrase, is only one example of the *descriptive fallacy*, so common in philosophy. Even if some language is now purely descriptive, language was not in origin so, and much of it is still not so. Utterance of obvious ritual phrases, in the appropriate circumstances, is not *describing* the action we are doing, but *doing* it ('I do'): in other cases it functions, like tone and expression, or again like punctuation and mood, as an intimation that we are employing language in some special way ('I warn', 'I ask', 'I define'). Such phrases cannot, strictly, *be* lies, though they can 'imply' lies, as 'I promise' implies that I fully intend, which may be untrue.

If these are the main and multifarious points that arise in familiar cases where we ask 'How do you know that this is a case of so-and-so?', they may be expected to arise likewise in cases where we say 'I know he is angry'. And if there are, as no doubt there are, special difficulties in this case, at least we can clear the ground a little of things which are not special difficulties, and get the matter in better perspective.

As a preliminary, it must be said that I shall only discuss the question of feelings and emotions, with special reference to anger. It seems likely that the cases where we know that another man thinks that 2 and 2 make 4, or that he is seeing a rat, and so on, are different in important respects from, though no doubt also similar to, the case of knowing that he is angry or hungry.

In the first place, we certainly do say sometimes that we know another man is angry, and we also distinguish these occasions from others on which we say only that we *believe* he is angry. For of course, we do not for a moment suppose that we *always* know, of *all* men, whether they are angry or not, or that we could discover it. There are many occasions when I

realize that I can't possibly tell what he's feeling: and there are many *types* of people, and many individuals too, with whom I (they being what they are, and I being what I am) never can tell. The feelings of royalty, for example, or fakirs or bushmen or Wykehamists or simple eccentrics—these may be very hard to divine: unless you have had a prolonged acquaintance with such persons, and some intimacy with them, you are not in any sort of position to know what their feelings are, especially if, for one reason or another, they can't or don't tell you. Or again, the feelings of some individual whom you have never met before—they might be almost anything: you don't know his character at all or his tastes, you have had no experience of his mannerisms, and so on. His feelings are elusive and personal: people differ so much. It is this sort of thing that leads to the situation where we say 'You never know' or 'You never can tell'.

In short, here even more than in the case of the goldfinch, a great deal depends on how familiar we have been in our past experience with this type of person, and indeed with this individual, in this type of situation. If we have no great familiarity, then we hesitate to say we know: indeed, we can't be expected to say (tell). On the other hand, if we *have* had the necessary experience, then we can, in favourable current circumstances, say we know: we certainly can recognize when some near relative of ours is angrier than we have ever seen him.

Further, we must have had experience also of the emotion or feeling concerned, in this case anger. In order to know what you are feeling, I must also apparently be able to imagine (guess, understand, appreciate) what you're feeling. It seems that more is demanded than that I shall have learned to discriminate displays of anger in others: I must also have been angry myself.[1] Or at any rate, if I have never felt a certain

[1] We say we don't know what it must feel like to be a king, whereas we do know what one of our friends must have felt when mortified. In this ordinary (imprecise and evidently not whole-hog) sense of 'knowing what it

emotion, say ambition, then I certainly feel an *extra* hesitation in saying that his motive is ambition. And this seems to be due to the very special nature (grammar, logic) of feelings, to the special way in which they are related to their occasions and manifestations, which requires further elucidation.

At first sight it may be tempting to follow Mr. Wisdom, and to draw a distinction between (1) the physical symptoms and (2) the feeling. So that when, in the current case, I am asked 'How can you tell he's angry?' I should answer 'From the physical symptoms', while if *he* is asked how *he* can tell he's angry, he should answer 'From the feeling'. But this seems to be a dangerous over-simplification.

In the first place, 'symptoms' (and also 'physical') is being used in a way different from ordinary usage, and one which proves to be misleading.

'Symptoms', a term transferred from medical usage,[1] tends to be used only, or primarily, in cases where that of which there are symptoms is something undesirable (of incipient disease rather than of returning health, of despair rather than of hope, of grief rather than of joy): and hence it is more colourful than 'signs' or 'indications'. This, however, is comparatively trivial. What is important is the fact that we never talk of 'symptoms' or 'signs' except *by way of implied contrast with inspection of the item itself.* No doubt it would often be awkward to have to say exactly where the signs or symptoms end and the item itself begins to appear: but such a division is always implied to exist. And hence the words 'symptom' and 'sign' have no use except in cases where the item, as in the case

would be like' we do often know what it would be like to be our neighbour drawing his sword, whereas we don't know (can't even guess or imagine), really, what it would feel like to be a cat or a cockroach. But of course we don't ever 'know' what in our neighbour accompanies the drawing of his sword in Mr. Wisdom's peculiar sense of 'know what' as equivalent to 'directly experience that which'.

[1] Doctors nowadays draw a distinction of their own between 'symptoms' and '(physical) signs': but the distinction is not here relevant, and perhaps not very clear.

of disease, is liable to be *hidden*, whether it be in the future, in the past, under the skin, or in some other more or less notorious casket: and when the item is itself before us, we no longer talk of signs and symptoms. When we talk of 'signs of a storm', we mean signs of an impending storm, or of a past storm, or of a storm beyond the horizon: we do *not* mean a storm on top of us.[1]

The words function like such words as 'traces' or 'clues'. Once you know the murderer, you don't get any more clues, only what were or would have been clues: nor is a confession, or an eye-witness's view of the crime, a particularly good clue—these are something different altogether. When the cheese is not to be found or seen, then there may be traces of it: but not when it's there in front of us (though of course, there aren't, then, 'no traces' of it either).

For this reason, it seems misleading to lump together, as a general practice, all the characteristic features of any casual item as 'signs' or 'symptoms' of it: though it is of course sometimes the case that some things which could in appropriate circumstances be called characteristics or effects or manifestations or parts or sequelae or what not of certain items may *also* be called signs or symptoms of those items in the appropriate circumstances. It seems to be this which is really wrong with Mr. Wisdom's paradox (Other Minds III) about looking in the larder and finding 'all the signs' of bread, when we see the loaf, touch it, taste it and so on. Doing these things is not finding (some) signs of bread at all: the taste or feel of bread is not a sign or symptom of bread at all. What I might be taken to mean if I announced that I had found signs of bread in the larder seems rather doubtful, since bread is not normally

[1] There are some, more complicated, cases like that of inflation, where the signs of incipient inflation are of the same nature as inflation itself, but of a less intensity or at a slower tempo. Here, especially, it is a matter for decision where the signs or 'tendencies' end and where the state itself sets in: moreover, with inflation, as with some diseases, we can in some contexts go on talking of signs or symptoms even when the item itself is quite fairly decidedly present, because it is such as not to be patent to simple observation.

casketed (or if in the bin, leaves no traces), and not being a transeunt event (impending bread, &c.), does not have any normally accepted 'signs': and signs, peculiar to the item, have to be more or less normally accepted. I might be taken to mean that I had found traces of bread, such as crumbs, or signs that bread had at one time been stored there, or something of the kind: but what I could *not* be taken to mean is that I had seen, tasted, or touched (something like) bread.

The sort of thing we do actually say, if the look is all right but we haven't yet tasted it, is 'Here is something that looks like bread'. If it turns out not to be bread after all, we might say 'It tasted like bread, but actually it was only bread-substitute', or 'It exhibited many of the characteristic features of bread, but differed in important respects: it was only a synthetic imitation'. That is, we don't use the words 'sign' or 'symptom' at all.

Now, if 'signs' and 'symptoms' have this restricted usage, it is evident that to say that we only get at the 'signs' or 'symptoms' of anything is to imply that we never get at *it* (and this goes for '*all* the signs' too). So that, if we say that I only get at the *symptoms* of his anger, that carries an important implication. But *is* this the way we do talk? Surely we do not consider that we are never aware of more than *symptoms* of anger in another man?

'Symptoms' or 'signs' of anger tend to mean signs of *rising* or of *suppressed* anger. Once the man has exploded, we talk of something different—of an expression or manifestation or display of anger, of an exhibition of temper, and so forth. A twitch of the eyebrow, pallor, a tremor in the voice, all these may be symptoms of anger: but a violent tirade or a blow in the face are not, they are the acts in which the anger is vented. 'Symptoms' of anger are not, at least normally, contrasted with the man's own inner personal feeling of anger, but rather with the actual display of anger. Normally at least, where we have only symptoms to go upon, we should say only that we *believe*

that the man is angry or getting angry: whereas when he has given himself away we say that we *know*.[1]

The word 'physical' also, as used by Mr. Wisdom in contrast to 'mental', seems to me abused, though I am not confident as to whether this abuse is misleading in the current case. He evidently does not wish to call a man's feelings, which he cites as a typical example of a 'mental' event, *physical*. Yet this is what we ordinarily often do. There are many physical feelings, such as giddiness, hunger or fatigue: and these are included by some doctors among the physical signs of various complaints. Most feelings we do not speak of as either mental or physical, especially emotions, such as jealousy or anger itself: we do not assign them to the *mind* but to the *heart*. Where we do describe a feeling as mental, it is because we are using a word normally used to describe a physical feeling in a special transferred sense, as when we talk about 'mental' discomfort or fatigue.

It is then, clear, that more is involved in being, for example, angry than simply showing the symptoms and feeling the feeling. For there is also the display or manifestation. And it is to be noted that the feeling is related in a unique sort of way to the display. When we are angry, we have an impulse, felt and/or acted on, to do actions of particular kinds, and, unless we suppress the anger, we do actually proceed to do them. There is a peculiar and intimate relationship between the emotion and the natural manner of venting it, with which, having been angry ourselves, we are acquainted. The ways in which anger is normally manifested are *natural* to anger just

[1] Sometimes, it is said, we use 'I know' where we should be prepared to substitute 'I believe', as when we say 'I know he's in, because his hat is in the hall': thus 'know' is used loosely for 'believe', so why should we suppose there is a fundamental difference between them? But the question is, what exactly do we mean by 'prepared to substitute' and 'loosely'? We are 'prepared to substitute' *believe* for *know* not as an *equivalent* expression but as a weaker and therefore preferable expression, in view of the seriousness with which, as has become apparent, the matter is to be treated: the presence of the hat, which would serve as a proof of it's owner's presence in many circumstances, could only through laxity be adduced as a proof in a court of law.

as there are tones *naturally* expressive of various emotions (indignation, &c.). There is not normally taken to be[1] such a thing as 'being angry' apart from any impulse, however vague, to vent the anger in the natural way.

Moreover, besides the natural expressions of anger, there are also the natural *occasions* of anger, of which we have also had experience, which are similarly connected in an intimate way with the 'being angry'. It would be as nonsensical to class these as 'causes' in some supposedly obvious and 'external' sense, as it would be to class the venting of anger as the 'effect' of the emotion in a supposedly obvious and 'external' sense. Equally it would be nonsensical to say that there are three wholly distinct phenomena, (1) cause or occasion, (2) feeling or emotion, and (3) effect or manifestation, which are related together 'by definition' as all necessary to anger, though this would perhaps be less misleading than the other.

It seems fair to say that 'being angry' is in many respects like 'having mumps'. It is a description of a whole pattern of events, including occasion, symptoms, feeling and manifestation, and possibly other factors besides. It is as silly to ask 'What, really, *is* the anger *itself*?' as to attempt to fine down 'the disease' to some one chosen item ('the functional disorder'). That the man himself feels something which we don't (in the sense that he feels angry and we don't) is, in the absence of Mr. Wisdom's variety of telepathy,[2] evident enough, and incidentally nothing to complain about as a 'predicament': but there is no call to say that 'that' ('the feeling')[3] *is* the *anger*. The pattern of events

[1] A new language is naturally necessary if we are to admit unconscious feelings, and feelings which express themselves in paradoxical manners, such as the psycho-analysts describe.

[2] There is, it seems to me, something which does actually happen, rather different from Mr. Wisdom's telepathy, which does sometimes contribute towards our knowledge of other people's feelings. We do talk, for example, of 'feeling another person's displeasure', and say, for example, 'his anger could be felt', and there seems to be something genuine about this. But the feeling we feel, though a genuine 'feeling', is *not*, in these cases, displeasure or anger, but a special *counterpart* feeling.

[3] The 'feelings', i.e. sensations, we can observe in ourselves when angry are

whatever its precise form, is, fairly clearly, peculiar to the case of 'feelings' (emotions)—it is not by any means exactly like the case of diseases: and it seems to be this peculiarity which makes us prone to say that, unless we have had experience of a feeling ourselves, we cannot know when someone else is experiencing it. Moreover, it is our confidence in the general pattern that makes us apt to say we 'know' another man is angry when we have only observed parts of the pattern: for the parts of the pattern are related to each other very much more intimately than, for example, newspapermen scurrying in Brighton are related to a fire in Fleet Street.[1]

The man himself, such is the overriding power of the pattern, will sometimes accept corrections from outsiders about his own emotions, i.e. about the correct description of them. He may be got to agree that he was not really angry so much as, rather, indignant or jealous, and even that he was not in pain, but only fancied he was. And this is not surprising, especially in view of the fact that he, like all of us, has primarily learnt to use the expression 'I am angry' of himself by (*a*) noting the occasion, symptoms, manifestation, &c., in cases where other persons say 'I am angry' of *themselves* (*b*) being told by others, who have noted all that can be observed about *him* on certain occasions, that 'You are angry', i.e. that he should say 'I am angry'. On the whole, 'mere' feelings or emotions, if there are such things genuinely detectable, are certainly very hard to be sure about, even harder than, say, tastes, which we already choose to describe, normally, only by their occasions (the taste 'of tar', 'of pineapple', &c.).

All words for emotions are, besides, on the vague side, in two ways, leading to further hesitations about whether we 'know' when he's angry. They tend to cover a rather wide and ill-defined variety of situations: and the patterns they cover

such things as a pounding of the heart or tensing of the muscles, which cannot in themselves be justifiably called 'the feeling of anger'.

[1] It is therefore misleading to ask 'How do I get from the scowl to the anger?'

tend to be, each of them, rather complex (though common and so not difficult to recognize, very often), so that it is easy for one of the more or less necessary features to be omitted, and thus to give rise to hesitation about what exactly we should say in such an unorthodox case. We realize, well enough, that the challenge to which we are exposed if we say we *know* is to *prove* it, and in this respect vagueness of terminology is a crippling handicap.

So far, enough has perhaps been said to show that most of the difficulties which stand in the way of our saying we know a thing is a goldfinch arise in rather greater strength in the case where we want to say we know another man is angry. But there is still a feeling, and I think a justified feeling, that there is a further and quite *special* difficulty in the latter case.

This difficulty seems to be of the sort that Mr. Wisdom raises at the very outset of his series of articles on 'Other Minds'. It is asked, might the man not exhibit all the symptoms (and display and everything else) of anger, even ad infinitum, and yet still *not* (*really*) *be* angry? It will be remembered that he there treats it, no doubt provisionally, as a difficulty similar to that which can arise concerning the reality of any 'material object'. But in fact, it has special features of its own.

There seem to be three distinguishable doubts which may arise:

1. When to all appearances angry, might he not really be labouring under some other emotion, in that, though he normally feels the same emotion as we should on occasions when we, in his position, should feel anger and in making displays such as we make when angry, in this particular case he is acting abnormally?

2. When to all appearances angry, might he not really be labouring under some other emotion, in that he normally feels, on occasions when we in his position should feel anger and when acting as we should act if we felt anger, some feeling which we, if we experienced it, should distinguish from anger?

3. When to all appearances angry, might he not really be feeling no emotion at all?

In everyday life, all these problems arise in special cases, and occasion genuine worry. We may worry (1) as to whether someone is *deceiving* us, by suppressing his emotions, or by feigning emotions which he does not feel: we may worry (2) as to whether we are *misunderstanding* someone (or he us), in wrongly supposing that he does 'feel like us', that he does share emotions like ours: or we may worry (3) as to whether some action of another person is really deliberate, or perhaps only involuntary or inadvertent in some manner or other. All three varieties of worry may arise, and often do, in connexion with the actions of persons whom we know very well.[1] Any or all of them may be at the bottom of the passage from Mrs. Woolf:[2] all work together in the feeling of loneliness which affects everybody at times.

None of these three special difficulties about 'reality' arises in connexion with goldfinches or bread, any more than the special difficulties about, for example, the oasis arise in connexion with the reality of another person's emotions. The goldfinch cannot be assumed, nor the bread suppressed: we may be deceived by the appearance of an oasis, or misinterpret the signs of the weather, but the oasis cannot lie to us and we cannot misunderstand the storm in the way we misunderstand the man.

Though the difficulties are special, the ways of dealing with them are, initially, similar to those employed in the case of the goldfinch. There are (more or less roughly) established procedures for dealing with suspected cases of deception or of misunderstanding or of inadvertence. By these means we do very often establish (though we do not expect *always* to establish) that someone is acting, or that we were misunder-

[1] There is, too, a special way in which we can doubt the 'reality' of our own emotions, can doubt whether we are not 'acting to ourselves'. Professional actors may reach a state where they never really know what their genuine feelings are.

[2] [Quoted by Wisdom in his contribution to this Symposium. Ed.]

standing him, or that he is simply impervious to a certain emotion, or that he was not acting voluntarily. These special cases where doubts arise and require resolving, are contrasted with the normal cases which hold the field[1] *unless* there is some special suggestion that deceit, &c., is involved, and deceit, moreover, of an intelligible kind in the circumstances, that is, of a kind that can be looked into because motive, &c., is specially suggested. There is no suggestion that I *never* know what other people's emotions are, nor yet that in particular cases I might be wrong for no special reason or in no special way.

Extraordinary cases of deceit, misunderstanding, &c. (which are themselves not the normal), do not, *ex vi termini*, ordinarily occur: we have a working knowledge of the occasions for, the temptations to, the practical limits of, and the normal types of deceit and misunderstanding. Nevertheless, they *may* occur, and there may be varieties which are common without our yet having become aware of the fact. If this happens, we are in a certain sense wrong, because our terminology is inadequate to the facts, and we shall have thenceforward to be more wary about saying we know, or shall have to revise our ideas and terminology. This we are constantly ready to do in a field so complex and baffling as that of the emotions.

There remains, however, one further special feature of the case, which also differentiates it radically from the goldfinch case. The goldfinch, the material object, is, as we insisted above, uninscribed and *mute*: but the man *speaks*. In the complex of occurrences which induces us to say we know another man is angry, the complex of symptoms, occasion, display, and the rest, a peculiar place is occupied by the man's own statement as to what his feelings are. In the usual case, we accept this statement without question, and we then say that we know (as it were 'at second-hand') what his feelings are: though of course 'at second-hand' here could not be used to imply that anybody but he could know 'at first-hand', and hence perhaps it is not

[1] 'You cannot fool all of the people all of the time' is 'analytic'.

in fact used. In unusual cases, where his statement conflicts with the description we should otherwise have been inclined to give of the case, we do not feel bound to accept it, though we always feel some uneasiness in rejecting it. If the man is an habitual liar or self-deceiver, or if there are patent reasons why he should be lying or deceiving himself on this occasion, then we feel reasonably happy: but if such a case occurred as the imagined one where a man, having given throughout life every appearance of holding a certain pointless belief, leaves behind a remark in his private diary to the effect that he never did believe it, then we probably should not know what to say.

I should like to make in conclusion some further remarks about this crucial matter of our believing what the man says about his own feelings. Although I know very well that I do not see my way clearly in this, I cannot help feeling sure that it is fundamental to the whole Predicament, and that it has not been given the attention it deserves, possibly just because it is so obvious.

The man's own statement is not (is not treated primarily as) a sign or symptom, although it can, secondarily and artificially, be treated as such. A unique place is reserved for it in the summary of the facts of the case. The question then is: 'Why believe him?'

There are answers that we can give to this question, which is here to be taken in the general sense of 'Why believe him ever?' not simply as 'Why believe him this time?' We may say that the man's statements on matters other than his own feelings have constantly been before us in the past, and have been regularly verified by our own observations of the facts he reported: so that we have in fact some basis for an induction about his general reliability. Or we may say that his behaviour is most simply 'explained' on the view that he does feel emotions like ours, just as psycho-analysts 'explain' erratic behaviour by analogy with normal behaviour when they use the terminology of 'unconscious desires'.

These answers are, however, dangerous and unhelpful. They are so obvious that they please nobody: while on the other hand they encourage the questioner to push his question to 'profounder' depths, encouraging us, in turn, to exaggerate these answers until they become distortions.

The question, pushed further, becomes a challenge to the very possibility of 'believing another man', in its ordinarily accepted sense, at all. What 'justification' is there for supposing that there is another mind communicating with you at all? How can you know what it would be like for another mind to feel anything, and so how can you understand it? It is then that we are tempted to say that we only mean by 'believing him' that we take certain vocal noises as signs of certain impending behaviour, and that 'other minds' are no more really real than unconscious desires.

This, however, is distortion. It seems, rather, that believing in other persons, in authority and testimony, is an essential part of the act of communicating, an act which we all constantly perform. It is as much an irreducible part of our experience as, say, giving promises, or playing competitive games, or even sensing coloured patches. We can state certain advantages of such performances, and we can elaborate rules of a kind for their 'rational' conduct (as the Law Courts and historians and psychologists work out the rules for accepting testimony). But there is no 'justification' for our doing them as such.

Final Note

One speaker at Manchester said roundly that the real crux of the matter remains still that 'I ought not to say that I know Tom is angry, because I don't introspect his feelings': and this no doubt is just what many people do boggle at. The gist of what I have been trying to bring out is simply:

1. *Of course* I *don't* introspect Tom's feelings (we should be in a pretty predicament if I did).
2. *Of course* I *do* sometimes know Tom is angry.

Hence

3. to suppose that the question 'How do I know that Tom is angry?' is meant to mean 'How do I introspect Tom's feelings?' (because, as we know, that's the sort of thing that knowing is or ought to be), is simply barking our way up the wrong gum tree.