

PUBLICITY, EXTERNALISM AND INNER STATES*

BARRY C. SMITH

The critic Cyril Connolly once pointed out that diarists don't make novelists. He went on to describe the problem for the would-be writer. "Writing for oneself: no public. Writing for others: no privacy" (Cyril Connolly, *Journal*).

Connolly's quip nicely illustrates a problem for the public and the private that persists when we turn to the public and private dimensions of mind and meaning: namely, that attempts to accommodate one are usually at the expense of the other.

The philosophical problem this paper addresses is how to reconcile the inner and conscious dimension of speech with its outer and public dimension. In speaking, we are consciously aware of the meanings of the words we use. You are aware of the meanings of the words you are reading now. And in speaking, you and I choose which words to use, aware it seems of what we are saying. In trying to describe some situation, we sometimes face choices as to which of two words would best express a feature of the situation, we sometimes choose between them because the meaning of the one word is more peculiarly apt to express what it is we are trying to say than the other. To be able to do choose words in this way, the meanings we attach to them have to be immediately available to us, within our conscious reach. If the meanings of words were not within our conscious reach, it is hard to see how we could find speech—our own or anyone else's—so much as intelligible.

From this stance as ordinary language users, our use and understanding of language is easy and unreflective. Employing the language or responding to its use by others seems effortless. It requires no special figuring out on our part, nor does it appear to be any sort of cognitive achievement to find words meaningful in the way we typically do. But if what a person means by his words is a matter of what he has in mind, of how things are with him, consciously-speaking, when he uses words or listens to them, how can the meanings of these words also be publicly accessible to others and serve in communication? How can our words have meanings that match what others take them to mean when they hear us utter those words?

Meaning must be publicly accessible to others for communication to succeed, and yet what is publicly accessible are facts about a speaker's linguistic behaviour. So how can what someone has in mind as he speaks—the meaning he attaches to his words—also be publicly accessible to others on the basis of linguistic behaviour alone? How can the meanings he gives his words match what other people take him to mean?

In using words, one does not measure one's use or understanding of a term against any public, or other, standard, one simply starts speaking, giving the words the meanings one simply takes them to have. But for communication to succeed one has to get one's meaning across to others, to make what one is saying available to them somehow. And one could not do so if the significance one attaches to one's words was a wholly private and personal affair. Somehow one must be using words with the meanings other people take them to have, even though in using them one does not consult others or measure one's understanding against theirs.

In order to meaning anything at all by my words, there must be more to their having the meanings they have for me than my simply taking them to mean whatever I take them to mean, and this surely suggests my taking words to mean something is answerable to something beyond my opinion. But does it involve conformity to a standard? Surely, there has to be a fact about what I mean for me to know what I mean. The

* Paper first given at a colloquium *Externalism-Internalism Dispute and the Theory of Interpretation* at the Institute of Philosophy of the Czech Academy of Sciences in Prague, April 9, 2002 and at *Sentence Meaning, Utterance Meaning/Text, Work* in Prague in 2003. Thanks to Bill Child and Guy Longworth for comments on an earlier draft, to Peter Pagin for discussion of publicity and to audiences at both conferences.

question is: how can the effortless and immediate impression of what I mean by the words I am using amount to knowledge—authoritative first-personal knowledge—of an objective fact about what I mean; a fact available to others on third-personal grounds? How does the existence of publicly and behaviourally available facts about what a speaker means by his words square with the speaker’s own immediate and effortless knowledge of what he means? As speakers, we do not have to observe our own linguistic output to see what significance we attach to expressions. We just know what our words mean without the benefit of evidence or inference. The facts of meaning, or of what a rule governing the correct use of a word requires of me, must be something I am immediately and intimately acquainted with. And yet the facts of meaning or the existence of rules (on which, according to the later Wittgenstein, the possibility of meaning depends) must be answerable to more than what I take them to be, and are not in any sense private objects of awareness.

How, then, is one to tell the story about meaning or rules of use, on the one hand, and our effortless and authoritative knowledge of them on the other? Is there a notion of linguistic meaning that can do justice both to the inner experience of comprehension and to the outer facts of language use?¹ Such an account would have to reconcile what is immediately available to speakers in their inner acts of comprehension with what is outwardly available to others in speakers’ public practice. Let us call this the Reconciliation Problem.

1 The Reconciliation Problem

More needs to be said about the precise nature of the Reconciliation Problem. To begin with, is it a problem in metaphysics or epistemology? It is hard at first to offer a definitive answer. The problem is how to reconcile the inner apprehension of what words mean with those meanings being publicly available to others through our linguistic acts. Were we to construe talk of “inner apprehension” as meaning no more than first-personal knowledge of meaning—the distinctively authoritative and immediate way we know what our words mean—and were we to treat the publicity of meaning as no more than the requirement that others can have third-personal knowledge of what we mean, the Reconciliation Problem would reduce to the epistemological issue of how we reconcile first- and third-personal knowledge of the same subject matter while preserving the fundamental asymmetry between the two kinds of knowledge. This is an important and difficult problem, to be sure, but the Reconciliation Problem goes further. We want an account of the metaphysics of meaning that makes it possible for meaning to be known in these distinctively different ways. Those who have addressed the publicity of meaning in the literature have mostly been concerned with metaphysical questions about the nature of linguistic meaning. Thinkers such as Wittgenstein, Quine, Davidson, Dummett and McDowell, have all sought to draw conclusions about what meaning must be like if it is to be publicly accessible. However, their metaphysical conclusions are often motivated by epistemological concerns: what must meaning be like in order to be *known* on the basis of observing people’s linguistic practice.² The Reconciliation Problem incorporates this problem but it raises an additional one too, requiring us to say what meaning must be like in order for speakers to have *immediate conscious awareness* of it? If we could provide a satisfactory metaphysics meeting all the epistemological demands, we would have solved the Reconciliation Problem. So any way to reconcile first- and third-person knowledge of meaning will have to show how meaning can be immediately accessible to the inner component of speaking and understanding, and yet outwardly discernible in observable facts about public practice.

The involvement of both epistemological and metaphysical strands in the Reconciliation Problem qualifies it as an instance of what Christopher Peacocke has called the Integration Challenge:

The Integration Challenge in its general form is that of reconciling our metaphysics of any given area with our epistemology for that same area (Peacocke 1998, 349).

¹ This way of stating the problem follows Jonathan Lear in his The disappearing “we” to which the setting up in the present paper owes a great deal. Lear’s paper presents one version of what I am calling here the reconciliation problem. It is a version he finds in the work of the later Wittgenstein concerning “the relation between my inner experience of comprehension and my objective ability to use a language” (Lear 1984, 224).

² I am heavily indebted here to Peter Pagin’s insightful discussion of publicity in his Publicness and indeterminacy (2000).

In this case, the challenge is to find a plausible metaphysics of meaning that respects the two sides to our epistemology of understanding. Does the solution require a substantial metaphysics given in terms of subjective and objective aspects of language use? That depends on whether the question about how the inner and the outer are related is a metaphysical question, or whether the notions of inner and outer are themselves merely reflections of the different epistemological access one can have to the same facts about meaning. This remains to be seen.

The issue of reconciling the outer and the inner aspect of speech is also apt to call to mind debates between *individualists* and *externalists* over whether meanings are in the head, or whether the meaning of words depends on external factors in the speaker's social or physical environment. However, the issues are not the same, for the Reconciliation Problem is fac

about anything. For Frege, private associations with a word cannot be communicated, and so cannot be part of its meaning. Frege equates meaning with what is communicable, what can be shared. Anything that cannot be shared or conveyed to another is simply no part of the linguistic meaning of an expression. In this spirit, Michael Dummett claims that no part of an expression's meaning can:

contain anything which is not manifest in the use made of it, laying solely in the mind of the individual who apprehends that meaning (Dummett 1978, 216).

And from here it is short step to the position both McDowell and Dummett subscribe to, that, as McDowell puts it:

the significance of others' utterances in a language must, in general, lie open to view, in publicly available facts about linguistic behaviour in its circumstances (McDowell 1981, 314).

How did we reach this position? We began with the idea of it being possible to know the meanings of other people's words, and advanced from there to a fairly strong reading of the publicity of meaning according to which the meaning of people's words must be located on the surface of their speech. Let us call this the *exteriorising* move⁴. The question to ask is whether denial of meaning-privacy requires this stronger reading of publicity; whether, that is, the exteriorising move is warranted. Notice, that if it is, the Reconciliation Problem becomes particularly acute: how can meaning be immediately present to the speaker's mind if it also lies open to view on the surface of linguistic behaviour? There are proposals that attempt to address this question. These will be considered below, but at this stage the question is whether we need to confront the problem in just this form. Minimally, all the denial of privacy requires is that it be possible to know what someone means by their words. It is a further question how such knowledge is arrived at. The constraint in (P) says nothing about what *makes* the meanings of someone's words available to others; *a fortiori* it does not say that it is due to meaning's being present on the surface of speech. So must someone who subscribes to (P) conclude that

... when we want to understand meaning and communication we should not turn inward, towards mental states, but outward, to what is publicly observable (Føllesdal 1990, 98).

There would be reason for endorsing the exteriorising move if failure to endorse it left one with one no alternative but the hopeless picture just canvassed; the picture according to which the significance of another's words is a matter of mere guesswork as to how things are in a private sphere. If the move to make meaning public by locating it on the surface of speech represented the only escape from this picture then we would be forced to locate linguistic significance in speech behaviour. But so far it has not been established that this thesis does represent the only escape from the psychologism of the hopeless picture.

Consider, for instance, Frege's anti-psychologism. Frege's demand that the meanings, or senses, of expressions be communicable or shareable—a demand amounting to a rejection of private meanings—does not lead him to adopt the view that meanings themselves are located on the surface of speech. For Frege, the publicity requirement that the very same thought can be grasped by different thinkers commits one to a conception of the meaning of sentences in which those thoughts are expressed as belonging to an abstract realm, standing over and against us all. Thus Platonism about meaning represents a radically different alternative to the hopeless picture; though it may be one we consider equally unpalatable for its lack of a credible epistemology of meaning, leaving us with a problem in both the first-personal and the third-personal account of knowing what words (Platonistically) mean. An argument could then be offered to the effect that if one rejects Platonism about linguistic meaning, then the only way to avoid a hopeless psychologism is to locate meanings on the surface of linguistic behaviour. But even here, there is no reason to think that this is the only other alternative to a Platonistic or psychologistic conception of meaning. I shall argue that there is a plausible alternative to the exteriorising move that has up until now been overlooked, and I will provide a sketch of what that plausible alternative might look like. In the meantime, let us consider more direct attempts to support the exteriorising move.

⁴ It would be more natural to say "externalising" here, but the suggestion of externalism is one I want to avoid since it is part of the purpose of this paper to explore the relations between the exteriorising move and externalism about mind and meaning. The contrasting notion of interiority is not, I shall argue, incompatible with externalism. The term "interiority" is McDowell's.

3. Motivating Publicity

As mentioned above, an important part of understanding the Reconciliation Problem is getting clear about the nature of the claims about the publicity of meaning. One way to understand it is as a claim about the limitations on the individual's ability to attach meaning to his or her words. According to (P) we cannot mean anything by our words unless others could know the meanings we attach to them. It is important not to lose sight of this key feature of publicity. It insists that meaning something is more than merely having the subjective impression that one has given a word some significance. Further conditions must be met for there to be an objective meaning attaching to what the speaker is attempting to say. Another way of understanding publicity is as a claim about what meaning must be like in order to be detectable from the observation of people's linguistic behaviour. This idea carries metaphysical implications about the nature of meaning. Meaning must either be displayed on the surface of linguistic behaviour or be in some way reducible to it for there to be publicly discernible meanings: for us to have access, that is, to the meanings of one another's utterances. Now this strong thesis is what I want to reject but in doing so I do not want to give up on the first idea behind publicity. For it expresses an important insight and re-affirms that our attempts at communication are not just matters of speculation or guesswork. It also insists that meaning something by one's words is not an essentially private matter. Publicity imposes an epistemological condition on meaning something by words, viz. the condition that it must be possible for others to know the meanings someone attaches to his or her words. A central aim of this paper is to retain this condition while jettisoning the metaphysical claim about the location of meaning on the surface of behaviour.

The epistemological condition is motivated in large part by the later Wittgenstein who at *Philosophical investigations* §43 offers us a conception of meaning as use. Wittgenstein's position succeeds in combining a rejection of Fregean Platonism with an adherence to Frege's anti-,

seldom argued—is all that can make up the facts about meaning. So the positing of meaning facts must locate them in what we display to others in our observable behaviour in those circumstances. Hence the best notion of meaning Quine can reconstruct from these meagre materials is that of a stimulus-meaning. This is a publicly manifestable notion of meaning for an expression elicited as a response to observable circumstances. The notion serves well for the behaviouristic study of what others mean by their words but it offers little or no account of the speaker's point of view. The speaker is simply disposed to produce an expression under certain stimulus conditions and not under other conditions. But what of the speaker's understanding? After all, we don't just use words: we understand them. We do not use them blindly, nor do we find out what we mean or whether we understand a word by seeing whether we are able to use it. Asked whether you understand the word "pleached" you don't say, "Well let me try using it and see how I do." It just isn't like this: we simply know whether we understand a word or not, and when we do, we know what they mean. And we know this in advance of any particular use we make of them.⁵ So what is Quine's view of understanding? Even here, Quine attempts to account for the speaker's understanding third-personally, and behaviourally, in terms of use: "... understanding a word consists in knowing how to use it in sentences and how to react to such sentences" (1990, 58). Behavioural manifestation may well be a test of understanding, but it cannot *constitute* our understanding: knowing what a word means is not just being able to use it (i.e. produce it when prompted by a stimulus). We typically know what we mean when we speak, and yet dispositions, like the dispositions to use words, are not introspectible. Do I have the disposition to jump a metre? I don't know. I would need to try and see what happens. But as we have seen it isn't like this with words. Of course we need not construe abilities to use words in terms of behavioural dispositions. Perhaps abilities have insides as well as outsides⁶. This seems like the only way to reconcile use and the understanding that prompts use, for it is hard to see how merely behaving in certain ways could amount to knowing what we mean by our words.⁷ The moral is this: *outer states stand in need of inner experiences*.⁸

Is the very move of securing knowledge of the significance of words by locating their significance in observable behaviour—*exteriorising move*—already a form of externalism? It is not altogether clear. Though it has been interpreted this way by Davidson:

Quine teaches us that what a speaker means by what he says, and hence the thoughts that can be expressed in language, are not accidentally connected with what a competent interpreter can make of them, and this is a powerful externalist thesis (Davidson 2001, 11).

Davidson, like Quine, subscribes to a strong publicity thesis, which requires the exteriorising move; and both thinkers take this to have consequences for the indeterminacy of meaning and reference.

The semantic features of language are public features. What no one can, in the nature of the case, figure out from the totality of the relevant evidence cannot be part of meaning (Davidson 1984, 235).

However, Davidson's version of strong publicity and the exteriorising move differs from Quine's. The semantic features of language that are displayed publicly are not reduced to behaviour: they are displayed

⁵ In What do I know when I know a language? (in his 1993, 94–105), Dummett offers the example of the society hostess in a P. G. Wodehouse novel who is asked whether she speaks Spanish and replies, "I don't know. I've never tried." I take it that the absurdity (and pretension) of the reply contrasts sharply with the normal case.

⁶ Dummett, of course, comes close in a number of places to saying that the behaviour in which meaning is manifested amounts to what it is to know the meaning of expressions. However, he also claims that speech is a conscious rational activity and only those regularities consciously selected count as part of the language. The thought that abilities may be cognitive abilities with "insides" as well as "outsides" is perhaps what Dummett is after when he talks of linguistic abilities as being more than merely practical abilities and as having an ineliminable theoretical component that guides as speaker as to which uses to make of his expressions (see Dummett 1983, 112).

⁷ Bernard Williams once reminded me about the danger of the "it's hard to see" form of argument in philosophy. One is always at the mercy of the reply, "Not as hard as you think!" But in this particular case, I hope he would think we are on firm ground.

⁸ It is interesting to note that despite his behaviourism, even Quine is not free of the tension between first and third person perspectives. There is a tension between Quine's behaviourist and exterior account of linguistic significance (extracted though radical translation) and the individual's foundationalist need to construct a theory of the world—a web of belief—from the sequences of stimulations of his sensory surfaces.

in bits of behaviour *interpreted in intentional terms*. All we can mean, for Davidson, is all we can be known to mean by a *fully informed interpreter*. The facts are open to view when the behaviour is re-described by an interpreter who imposes intentional descriptions on the otherwise physical facts. Interpretation consists in the application of irreducible intentional and semantical concepts—concepts of mind and meaning—to the observable facts about the speaker's behaviour. The semantical and intentional notions are not reducible to behaviour but they do supervene on behaviour. Interpretation nets all the facts there are about mind and meaning: what a speaker means or believes is what he can be justifiably interpreted as meaning or believing when a theory is brought to bear to make sense of his behaviour in rational terms. Davidson's epistemology of mind and meaning are resolutely public and third-personal, but the issue which remains to be addressed is how the subject can know the facts about his mind or meaning when he does not interpret himself. How does he know these facts, about which he is authoritative, without recourse to the resources of an interpreter?

Davidson needs an account of the first-person knowledge a speaker or thinker enjoys about his own beliefs and meanings. The need to accommodate the first-person is made clear by the asymmetries between how the subject knows what he means or what he is thinking and how others know. They rely on evidence and he does not. The subject is authoritative about what he means or thinks, they are not.

Is there an inevitable tension between self-knowledge and the publicity of mind and meaning? It might appear that the public side to mind and meaning threatens the possibility of self-knowledge. But appearances to the contrary, the external dimension ensures that whether a subject is in a given mental state, or means such and such by his or her words, is an objective, or, at any rate, inter-subjective, matter, answerable to more than just the subject's opinions. The possibility of third-person knowledge serves as a *requirement* on a satisfactory account of self-knowledge. It ensures that our inner world does not stand alone, removed from the rest of reality.

Thus states that are immediately available to us as part of our inner world can be credited with an objective reality over and above our immediate impression of them. The hard problem is explaining how one's immediate impressions (of how things seem to one) can amount to knowledge of objective, empirical facts—one's being in certain publicly determinable states of mind, or one's uttering words with publicly interpretable meanings.

The problem arises when we try to square this objective and outward aspect of the mental with the special way in which we know our own minds from the first-person perspective. How can states whose natures belong partly in the public sphere be so readily available to us from the first-person point of view? (Smith 1998, 392)

Pace Davidson and Quine, we are inclined to think of what we mean and what we think as being automatically available to us as part of our inner lives. But now we see the other half of the problem:

How can the facts of mental life be part of the inner world of a subject if they are also objective facts knowable by others on the basis of outwardly observable behaviour? (*ibid.*)

The existence of a third-person perspective, that of an informed interpreter, secures a certain objectivity for claims about the details of a person's psychology. But what a fully satisfying account of psychological self-knowledge that respects Davidson's strong publicity thesis would have to explain is why voiced opinions (avowals) about our own psychological states, not arrived at by means of interpretation, are nonetheless answerable, for their correctness, to how we can be interpreted by others (i.e. to the pronouncements of an interpretation theory).

There are, of course, deflationary treatments of our psychological self-knowledge that try to account for the asymmetries between the first- and third- person perspectives without crediting subjects with an interior standpoint. Deflationary accounts try to acknowledge the first-person point of view by honouring a *mark* of first-person knowledge: namely, first-person authority, construed as the presumption of correctness attaching to psychological self-opinions not based on evidence or inference. This is an important aspect of the first person present tense psychological judgements but it is nowhere near enough to capture the first-person perspective of the agent. Accounts by Crispin Wright and Donald Davidson seek to save the presumption of correctness for self-opinions but offer no more than a deflationary account of linguistic and psychological self-knowledge. It is hard to recognise the phenomena they present as mental states of a subject's knowing what he means or believes, and so there is no real accommodation of first-person perspective.

4 Deflationary Accounts of the First-Person Perspective

According to Wright, meaning “cannot lie within the province of individual psychology” on pain of failing to be communicable to others. Moreover, if there is nothing other than a subject’s self-opinions for judgements about the application of a word to answer to then the object of such opinions fails to stabilize as a form of objective meaning. For these reasons, Wright locates the meaning of a term in patterns of use conforming to rules enshrined in the public practice. In this way, meaning is fully manifest in publicly observable patterns, leaving little or no room for the individual’s immediate awareness of what he means when using words as he does. What is it, then, for the speaker to know what he is saying, to use a word deliberately in accordance to its meaning, or to adhere to a particular rule for use? It cannot be that he waits for his use of words to reveal patterns sufficient for his interpretation of his own speech. That is of dubious coherence. Wright tells us that “there is no *essential* inner epistemology of rule-following” (Wright 2001, 188). Instead the subject may simply make pronouncements about what he judges to be the correct use of a word on a particular occasion. What makes such judgements correct?

How can judgements lack a substantial epistemology in this way, and yet still be *objective*—still have to answer to something distinct from our actual dispositions of judgement? (Wright 2001, 191)

The deflationary answer is that such self-pronouncements are true by default and stay true so long as no interpretation of the subject makes better sense of his words and deeds by overruling or discounting what he says about himself. The subject’s self-pronouncements must chime in with the ascriptions of attitudes and meanings the interpreter is prepared to make in order to give intelligible sense to his acts and utterances.⁹ The subject’s self-pronouncements are credited with pride of place in an account that makes interpretative sense of his actions and utterances. They will be true by default, unless they have to be overturned by an overall account of his behaviour that makes better sense of him in rational terms by discounting or overruling those opinions.¹⁰

This essentially third-personal account of the truth of what the subject says about what he means or thinks—what the subject says is true when it accords with what the interpreter would say—ensures first-person authority, the presumption of correctness of present-tense psychological self-ascriptions, but it does little to explain why these interpretively-established correct pronouncements count as *the subject’s knowing* what he means or thinks.

Knowing one’s own intentions or meanings “is not really a matter of ‘access to’—being in cognitive touch with—a state of affairs at all.” And the

authority standardly granted to a subject’s own beliefs, or expressed avowals about his intentional states [and meanings] is a constitutive principle: something which is not a consequence of the nature of those states [or meanings] and of an associated epistemologically privileged relation in which the subject stands to them, but enters primitively into the conditions for identification of what a subject believes, ... intends [and means] (Wright 2001, 312).

We simply issue self-pronouncements true by the interpreter’s lights. This offers little illumination of the first-person point of view, despite Wright’s claim that “phenomenologically, at any rate, construal of a novel utterance is often immediate and spontaneous” (Wright 1989, 190). The deflationary account of the first-person point of view is not so much a solution to the Reconciliation Problem as a way to avoid it by downplaying of the two perspectives we are seeking to reconcile.

Exclusive focus on the perspective of the interpreter neglects, at cost, the phenomenology of thought and meaning. However, restoring it requires more than merely adding an extraneous accompaniment to the external manifestation of meaning. The speaker must be accredited with *knowledge* of what he means or thinks, knowledge that he acquires in a way that is not available to others since *he* knows what he means or

⁹ See Wright (2001, 314), where he identifies “the correctness conditions of ascriptions” with “their interpretative utility.”

¹⁰ “... we do not *cognitively interact* with states of affairs which confer truth upon our opinions concerning our own intentional states [or meanings]; rather, we are, as it were, inundated, day by day, with opinions [which we are ceaselessly but subcognitively moved to make] concerning our own intentional states [or the meanings of our utterances] for which truth is the default position, so to say. They count as true provided that we hold them and no good purpose is served, in another’s quest to find us intelligible, by rejecting them” (*ibid.*, 313).

thinks without recourse to the external evidence by reference to which others are able to understand or interpret him. The problem that looms large is just how to explain his having knowledge of an objective fact about himself—a fact publicly accessible to others—in an effortless, non-evidential way. What makes the publicly accessible objective fact about what he means or what mental state he is in available to him immediately and effortlessly? It is clearly more than just a subjective matter of how things are in his consciousness. The Cartesian would have no difficulty in explaining how the subject could be so easily apprised of these facts since the facts—if such they are—enjoy no independence or distance from the subject himself. The nearness and lack of independence of the subject matter from the subject makes it easy to see why those matters should be so immediately available to him—what they cannot explain is how they could in any sense be available or accessible to others, since they have no existence independent of, and are entirely dependent on, the mind to which they belong. But if our rejection of Cartesianism and the essential subjectivism it entails, and our subsequent endorsement of publicity conditions as bringing into view genuine facts about meaning, is to leave room for the subject to know what he thinks and what he means—and thereby to know, albeit by entirely different means, exactly what it is that others who understand or interpret him know—we need more than just a conscious or phenomenological accompaniment to someone's issuing a meaningful (and thus interpretable) utterance or to his simply being in a mental state. The question therefore is how a subject's inner experience or awareness can amount to his knowledge of an objective fact about the significance of his thought or talk.

That question is just one side of the problem. The other side comes into view once we recognise that an agent's meaning or thinking something is bound up with his subjective experiences, his awareness of meaning and thinking. We are required to see how part of the agent's inner life could be available to others on the basis of his observable behaviour. Reconciling these two perspectives on the same phenomenon is what we are after, and a non-Cartesian externalism helps, rather than hinders, I shall argue.

A subtly different construal of the publicity of meaning is offered by John McDowell. He points out that:

Our attention is drawn to ... something present in the words—something capable of being heard or seen in the words by those who understand the language (1987, 99).

command of a language equips us to know one another's meaning without needing to arrive at that knowledge by interpretation, because it equips us to hear someone else's meaning in his words (1984, 350–1).

Remember, it can do so because:

the significance of others' utterances in a language must, in general, lie open to view, in publicly available facts about linguistic behaviour in its circumstances (1981, 314).

For McDowell, what we recognise in speech is presented on the surface of behaviour. This ensures that meaning cannot transcend use; a constraint necessary for a credible epistemology of understanding. Without it, according to McDowell, arriving at a view of what others mean would be a mere matter of guesswork as to “how things are in a private sphere concealed behind their behaviour” (*ibid.*)—a view opposed by Frege and Wittgenstein; the former because it threatened communicability and the objectivity of meaning; the latter because it flirted with an incoherent notion of our inner lives as insulated from the outer world.

Publicity, for McDowell (and Dummett, rather differently), is mandated as the only way to resist psychologism. We must locate linguistic significance on the surface of behaviour, i.e. take the exteriorising move, to resist the conception of meaning as residing in a private sphere. But what is available in the public sphere? When we speak all we can show others are bits of behaviour; we make noises and gestures, and from this meagre evidence they are somehow able to know what we are saying. If what is said must be gleaned from behaviour, where this is not a matter of guesswork or hypothesis, and not a matter of imposing an interpretation of that behaviour by means of theory, the significance of speech must be found *in* the behaviour itself. And yet if we follow Quine, this move leads to a reduction of meaning to stimulus meaning, which fails to comport with the richness of our linguistic understanding. McDowell wants to locate linguistic meaning on the surface of linguistic behaviour, but unlike Quine, McDowell refuses to limit descriptions of that behaviour to physical description. He suggests that the attempt to characterise linguistic behaviour in purely physical or bodily terms will leave us with no way to recognise the activity as

linguistic. Instead, we must describe the meaning-revealing behaviour in terms of those very meanings themselves. We cannot characterise it in anything less than linguistic terms, and still capture what is thereby revealed. Nor would we do justice to the phenomenology of understanding were we to describe speech in merely behavioural or physical terms. This is an important insight by McDowell, that we hear more in the speech sounds of a language that we understand than in one we don't. In a language that we understand we are immediately aware of what someone else is saying and cannot hear it merely as sound standing in need of interpretation. Thus, McDowell concludes, linguistic behaviour reveals meanings to be found there—for those with eyes to see (or ears to hear)—of the surface of linguistic practice. Meaning is publicly displayed in behaviour, but in behaviour richly construed as the saying that such and such. The meanings are there in observable behaviour but perceptible only by those who know the language. When we know the language there is no barrier to hearing the intelligibility of another's utterance, or hearing the meanings of their words: the meanings are all there on the surface of linguistic behaviour provided that we do not characterize the surface of speech in unnecessarily impoverished terms. Thus McDowell seems to have found a way to reconcile two competing perspectives: the behavioural means of transmitting meaning and the immediate phenomenological character of understanding. Understanding is a matter of direct perception because the meanings we perceive are out there on the surface of speech.

This externalises meaning without the reduction of content to the mere constructs out of behaviour, physically described, that Quine goes in for. (The epistemology is straightforward for McDowell because the metaphysics is rather extravagant.) McDowell, will tolerate a world of meaningful events. Meaning is out there for those with eyes to see or ears to hear: present in the speaker's words and as such there to be fielded by the attentive perceiver. But not just any perceiver—for as McDowell acknowledges, it is only those with knowledge of the language who are in a position to perceive the meaning facts. But then what does linguistic knowledge contribute to the situation and in what way does it help to make the meaning facts available? McDowell is somewhat reticent about the role of knowledge of language here. We get very little for instance on how it is acquired, what it consists in, and how it is exercised. For McDowell, linguistic understanding is simply a perceptual capacity—a capacity to perceive people's meaningful speech. But how can the linguistic knowledge that confers understanding be acquired? McDowell talks of training in a behavioural repertoire somehow bringing us to recognize a realm of meaning facts out there on the surface of behaviour: "How can drilling in a behaviour repertoire stretch one's perceptual capacities—cause one to be aware of facts of which one would not otherwise have been aware?" (*ibid.*, 333) Unsurprisingly, McDowell doesn't have an answer to this question, although he badly needs one. The suspicion being that we won't get an answer if we are asking the wrong question. But leaving aside possible answers, if possible, where do matters stand?

Does McDowell's bold picture give us everything we want? It depends how we should understand his important phenomenological insight. It is true that we *hear* the surface of speech (in a language we understand) as meaningful. This is not in doubt. But what enables us to hear meaning in people's words? And why do we hear more in a language we understand? McDowell takes the phenomenological claim to licence an epistemological one. The immediate experience of hearing meaning is a matter of directly perceiving the meanings that lie out there on the surface of behaviour. It is this metaphysical part of the story that is more difficult to make sense of. The epistemology is simple so long as we can tolerate the metaphysical extravagance of meanings being *out there* in the world, and believe that one comes to perceive more of the world by being inculcated into the behavioural practices of the speech community. McDowell insists that the world of meanings need not seem strange or magical once we see the world as re-enriched. But the anxiety about the magical picture is not relieved by replacing the word "magical" by "enriched." Rejecting the metaphysics need not mean rejecting the phenomenological insight or embracing psychologism. An adequate solution to the Reconciliation Problem must save McDowell's intuitions about the phenomenology but reject the epistemology and metaphysics. But first we must assemble the remaining materials for a satisfactory solution.

5 Externalism about Mind and Meaning

How do these issues play when we move to debates about externalist theses of meaning and content? Let us take externalism about the meaning of a person's words, or about the content of certain of their mental states, to commit one to the claim that the meaning of those words, or the content of those states, depends on factors external to the person's body, factors involving either the physical or the social environment. There may seem to be two threats to reconciliation posed by externalist doctrines. The first concerns a

supposed tension between externalism and self-knowledge. But as I suggested above this is not a real obstacle; in fact, the external perspective may provide an objective constraint on any candidate object of knowledge targeted by our self-opinions. The second threat, often pressed by individualists about the mental, is that externalist doctrines neglect crucial psychological considerations. This is a point worth attending to.

What are the crucial psychological considerations? Thought-experimental intuitions run by individualists often speak of how things seem from the inside when things change in the external or social environment. Should these intuitions always be respected? Not necessarily. Changes can occur without our noticing how things now differ and failure to admit this may lead to quite counter-intuitive results. Does it matter whether I am perceiving a friend or lover, or will any twin indistinguishable do as far as my thinking goes? Surely not. To give up on the ultimate authority of how things seem to the subject “from the inside” is not to give up on the inner altogether. Several externalists have placed emphasis on the importance of the phenomenological, and as we have seen McDowell seeks to support his externalism with an argument based on taking the phenomenology at face value as an encounter with the meanings in people’s words.¹¹

The real point is that it is the exclusive focus on the interpreter’s perspective, and not externalism per se, that leads to neglect of the phenomenology of understanding. The difficulty is that restoring it requires more than merely adding an extraneous accompaniment to the meaningful use of words. The experience of understanding must also amount to knowledge of what we mean. But is there a way to ensure this without taking the McDowellian route? If so, how can the phenomenological datum be accommodated while: (1) meeting the publicity constraint; (2) resisting the exteriorising move; (3) acknowledging that facts about meaning depend on factors external to the speaker?

Can meaning depend on external factors and yet be a phenomenological notion—part of the conscious experience of the speaker? This needs us to extend the mind and make minds accessible to others on the basis of their intentional behaviour. How so?

[We] should maintain that meanings, our own as well as others, can figure as integral components of our conscious life (McCulloch 1998, 265).

Moreover, these are indeed public features of the use of language, since the point is that public matters like speech, interpretation and communication are themselves fundamentally conscious phenomena. What I think is often what I put into words; and suitable audiences frequently hear *that same thing* (*ibid.*).

Meaning figures in phenomenology not only in the first-person case, where one is speaking consciously, but also in the third-person case when someone is speaking and one is aware of what they are saying ... It figures there directly, unmediated, but there is no need to claim that we are infallible detectors of meaning, our own or those of others: we can certainly misinterpret (hear the wrong meaning) ... (*ibid.*)

For McCulloch, meaning is a phenomenological notion. Meanings are manifest, and in being a phenomenological notion, meaning is what is directly heard or seen in the linguistic behaviour of speakers. But can we directly observe aspects of the inner lives of others? Are conscious phenomena directly observable features of the world? This is one problem. The other is to see how the mind could take in the objective features of the world on which the contents of these conscious experiences depend.

McDowell accepts that: “Knowledge of meanings is wholly a matter of how things are in a subject’s mind” (1992, 282). And “that a speaker means what she does by ‘water’ must be constituted at least in part by her physical and social environment.” This is not just an externalist thesis about meanings:

command of a word’s meaning is a mental capacity ... the mind [is] the locus of our manipulations of meanings ... Meanings are in the mind but as [Putnam’s] argument establishes, they cannot be in the head; therefore, we ought to conclude, the mind is not in the head” (McDowell 1992, 276).

¹¹ Psychological considerations, include the phenomenological insights, about the nature of linguistic understanding have been stressed by, among others, John McDowell (1998a; 1998b); Tyler Burge (1999); Gregory McCulloch (1998); Barry C. Smith (1992; 1998). There is also Wittgenstein’s “reminder” about meaning being all there at an instant. On a strictly phonological point about understanding Galen Strawson (1996) has suggested the need to acknowledge the experience of meaning, while Jerry Fodor (1983) has argued for a language module on the grounds of the mandatory operation of processors that ensure we hear words as meaningful and not merely as sounds.

So recognising other's meanings can be recognising what they have in mind even though that need not be hidden from us in some private or internal sphere. The mind is on show in behaviour, spread on the world. Meanings as parts of the mind are literally out there. McDowell supposes that meaning is available to view on the surface of linguistic practice because the surface of perceived speech is contentful and plays a significant part in the phenomenology of the hearer. Hence the behavioural surface of speech is a locus of meaning and significance, a place where meaning and mind reside.

6 An Alternative Solution

Instead of talking this way, and taking meanings to be "out there" on the surface of linguistic behaviour, I want to offer an account in which meanings occur within individual minds; but I also want to show why such an account need not commit us to seeing the significance individuals attach to their words as being either subjective or inaccessible to others.

To say meanings occur in individual minds is not to say they are private or subjective. What we each come to have in mind are the meanings we attach to a given speech sound, and this is fixed by a process of learning which essentially involves others. This process depends essentially upon the way we correlate our experiences with others in a shared perceptual environment. (More about this in a moment.) The link, in the early learning of language, between understanding and experience need not commit us to an account of meaning based on private states, where meaning is strictly inaccessible to others. Meaning something by one's words can involve experience without it rendering those meanings incommunicable. It may be thought that meanings are to be known about only by operating in accordance with Burke's Assumption: a position advanced by Edward Craig. According to Craig, meaning can involve private states, but we can know each others' meanings if we assume other people's private states are the same as ours in the similar circumstances. Craig offers this view of Edmund Burke's point:

... human beings naturally assume that others, when in broadly similar circumstances, experience inner states similar to those which they themselves experience (Craig 1997, 700).

Should we accept this view of individuals assuming their private states are the same? I think it will not fully serve as an account of the epistemology and phenomenology of understanding—in particular it's peculiar immediacy—but that is not to say that Burke's Assumption is wholly wrong. It contains a key insight we can build on. The insight is the ease with which we rely on the fact that others are in similar experiential states to ours, and we do so without need of inference. It is automatic in some cases to make the assumption, but this is only the default case. We need a way of connecting up the automatic assumption of similar experience with other facts about speakers and hearers in a way that explains the use of this information for shared knowledge of minds. In effect, we need to provide more detail about the starting point for Burke's Assumption and the way the default case provides grounds for knowledge. I shall set out the account I favour via an analogy with looking and seeing.

7 The Analogy with "Looking" and "Seeing"

When there is a looking there is usually a seeing: an individual looking in a certain direction typically sees what is going on there. Now, we can observe what others are looking at but we are not privy to their *seeings*. In our own case, we do not observe our *lookings*: we just see. We look in a given direction and see what is there. And in a relatively unproblematic sense we take ourselves to know what other people have seen by looking where they are looking and taking what we see to show us what they see. For example, if your attention is suddenly caught by something over my shoulder and I want to know what you've seen, I simply turn around, following the direction of your eye gaze, and by looking in that direction I now I see what you saw. My confidence that I know what you've seen is not based on an inference. I simply read off your state from mine. It is automatically assumed that if I look where you're looking I see what you see. We make this assumption all the time and are highly successful in doing so. What is more, these connections between lookings and seeings are established very early in childhood and give rise to no special sceptical problems about other minds, until, that is, we begin to do some philosophy. The connection between looking and seeing is so basic and well-established that we do not need to support it with any additional considerations. In this way, without inference, we automatically read off someone else's seeing from our seeings when we look where they are looking. This is the default case: in relatively

simple situations where we share broadly similar perspectives nothing leads us into error. Later in development we learn to adjust for the relative perspectives of ourselves and others, and the obtrusion of various objects in another's line of sight. Nevertheless, with these connections intact the fact that we can observe a looking but not a seeing, doesn't prevent us from moving easily from knowledge of one to knowledge of the other (at least in a sense of knowing what you've seen that matters). The account depends on our reading what we experientially see when looking in a given direction for what others see when looking in the same direction. For instance, when we see someone's attention is caught by something, we wonder what they've seen, and by retraining our gaze and looking where they are looking, we come to see something, and take what we are seeing to be what they are seeing—what they are attending to at that moment.

I want to suggest we can use the account of looking and seeing as an analogy for the connection in language between use and understanding. Just as in the visual case, we can observe people looking but are not privy to their seeing, so in the linguistic case, we can observe their use but are not privy to their experience of what they say. From the speaker's perspective, he doesn't observe his use of words, he understands what he says in using them. Now in the analogy I am envisaging, where there is a use there is understanding; and if I hear people using familiar words on a certain occasion, then I take what I understand by those words to be what you or anyone who uses them understands by them.

The insight to build on for an account of shared knowledge of word meaning is the way we move effortlessly from use to understanding, treating these notions as being wedded to one another the way we do with the close (by our lights, inseparable) connection between looking and seeing.

In learning words from others our experience comes to be coordinated with that of another person, and so comes to have a content conditioned by the contents of another's mind. The correlation between looking and seeing across individuals, the coordination of joint attention and the correlation between understanding and use all play an essential role here.

What I mean by a word in my vocabulary depends on how, and from whom, I learned the word. My experience is not wholly subjectively constituted or exhausted. It is what *I* mean by words—the meaning *I* attach to these expressions—that constitutes the significance words have for me. The first-person case is basic. And in the default case, I am entitled to take what I mean by the word to be what anyone who uses this word means by it. In this way, there is no difference for me between what I mean by a word and what the word means. (Of course this is open to defeat in non-basic cases where there is reason to suppose that others do not mean what I mean and resort to interpretation is necessary.) How have I succeeded in attaching meaning to words in my vocabulary; to the words I recognize? This is not a matter of finding out what you or others mean, or taking possession of an independently constituted meaning the words has. We can take advantage here of a key insight from Davidson:

Someone who is consciously teaching a beginner the use of a word may think of herself as simply passing on a meaning that already attaches the word. But from the learner's point of view, the word—the sound—is being *endowed* with a meaning (Davidson 2001, 14).

The insight is that for each speaker it is a matter of him or her investing the item heard with meaning and not simply recognising a meaning that is already there, let alone there on the surface of other people's meaningful speech. The coordination of the minds of speaker and hearer, teacher and learner, is what connects the minds of each and ensures that the meaning the hearer invests sounds with in his or her individual mind is tightly connected to, and in that sense utterly dependent on what is in the mind of the speaker or teacher (and so of others too). Coming to mean something is not just a matter of individuals acting in isolation and happening, correctly as it were, to assume others have those same internal states of significance-investing experiences. The picture offered by Burke's Assumption does nothing to suggest why speakers in the same environment will come to mean the same and attach just the meanings which would enable them to speak and share their thoughts with one another.

8 Understanding Others

Early on babies learn to coordinate their experience with others in the case of joint attention to objects and shared gaze with the mother. Later they learn to correct for perspective and to attribute seeings different from their own to other sighted beings. In the linguistic case, where there is a use there is an understanding, and just like in vision, we don't observe our own use, we just speak and have an understanding of what we

say. With others we never have their experience of understanding but we observe their use and we rely on the connection between use and understanding in our own case to take use and (our) understanding to go together to have knowledge of them. This model needs more elaboration but that is a task for another paper. I hope enough has been said here to sketch the shape such an account would take, basing learning of word meaning on joint attention. The notion of perceptual experience, the experience of the child is an externalist notion and the resulting coordination of minds brought about by these correlations make the resultant state essentially social though a matter of how things are with the individual/idiolectal language. Externalist claim about perceptual experience is a claim that perceptual experience has to be an encountering of objects, and an awareness of others. Needless to say this is a point McDowell is fond of making:

... since there is no rationally satisfactory route from experiences, conceived as, in general, less than encounters with objects, glimpses of objective reality, to the epistemic position we are manifestly in, experiences must be intrinsically encounters with objects (McDowell 1994, 193).

Meaning is secured and established when others share a common object of attention they are jointly attending to, and where our joint experience is commemorated by the sound introduced by the learner or teacher. The child comes to invest that word with meaning, it serves as a sound label for the object, and this securing of meaning depends essentially on others. However, the meaning the child hears and understands for any use of that word is, in the basic case, the meaning *he* invests the word with and not the meaning he “hears” others as enunciating. He simply takes that word to mean such and such and, if the conditions for acquisition of word meaning were met and he has not strayed beyond that immediate linguistic circle, he is right to do so. At the same time his default assumption, which can eventually be overturned, is that this is how anyone who uses the word understands it. His meaning is *the* meaning of the word. He takes the word to mean such and such—and within the sub-community of language users, he is right to do so. His experience of meaning is an experience in him, not an experience of something out there on the surface of behaviour, but it is an experience fixed in concert with others and the commonly experienced features of the environment and thus it is an externalist notion. He has no difficulty knowing what others are saying to him—in good cases—but this is because he knows the meaning the words have for him, the meaning he attaches to those sounds, and this he has no difficulty knowing. There can be an objective fact about both speaker and hearer meaning the same but not explained the way McDowell or McCulloch see the epistemology. There is the same meaning, the symmetry in truth conditions speaker and hearer attach to the sentence uttered, but there remains an epistemological asymmetry: I know what I mean just by taking myself to mean such and such. I take these words to mean such and such and I am usually right to do so. In the case of others, taking them to mean such and such by the default assumption that they will mean what I mean (or, by my lights, what anyone means) by these words carries no guarantee of correctness. There will be reliability but not the same grounds of entitlement. There is a riskiness in the case of others not there in the subject’s own case. So the epistemic symmetry remains between first-person and third-person perspective even though no phenomenological clue to this is given.¹²

The Reconciliation Problem, though not solved has now been properly addressed, and may be tackled in this new way. How should we construe publicity? Not through the exteriorising move. Meaning can be (tied to) use, but it also crucially depends on understanding. The publicity condition can be satisfied though not in the way we thought it could be. We respect i

School of Philosophy
 Birkbeck College
 University of London
 Malet Street
 London WC1E 7HX
 UK
 e-mail: b.smith@philosophy.bbk.ac.uk

References

- Burge, Tyler. 1999. Comprehension and interpretation. In *The philosophy of Donald Davidson*, edited by Lewis Edwin Hahn. Chicago: Open Court, 229–50.
- Craig, Edward. 1997. Meaning and privacy. In *A companion to the philosophy of language*, edited by Bob Hale and Crispin Wright. Oxford: Blackwell, 127–45.
- Davidson, Donald. 1984. *Inquiries into truth and interpretation*. Oxford: Clarendon Press.
- . 2001. Externalisms. In *Interpreting Davidson*, edited by Petr Kofárko, Peter Pagin and Gabriel Segal. Stanford: CSLI, 1–16.
- Dummett, Michael. 1978. *Truth and other enigmas*. London: Duckworth.
- . 1983. Language and truth. In *Approaches to language*, edited by Roy Harris. Oxford: Pergamon Press, 95–125.
- . 1993. *The seas of language*. Oxford: Oxford University Press.
- Fodor, Jerry. 1983. *The modularity of mind*. Cambridge, MA: MIT Press.
- Føllesdal, Dagfinn. 1990. Indeterminacy and mental states. In *Perspectives on Quine*, edited by Robert Barrett and Roger Gibson. Oxford: Blackwell, 98–109.
- Lear, Jonathan. 1984. The disappearing “we”. *Proceedings of the Aristotelian Society, Supplementary Volume* 58, 219–42.
- McCulloch, Gregory. 1998. Intentionality and interpretation. In *Current issues in philosophy of mind*, edited by Anthony O’Hear. Cambridge: Cambridge University Press, 253–71.
- McDowell, John. 1981. Anti-realism and the epistemology of understanding. Reprinted in McDowell (1998b), 314–43.
- . 1984. Wittgenstein on following a rule. *Synthese* 58:325–63.
- . 1987. In defence of modesty. Reprinted in McDowell (1998b), 87–107.
- . 1992. Putnam on mind and meaning. Reprinted in McDowell (1998b), 275–91.
- . 1994. The content of perceptual experience. *The Philosophical Quarterly* 44:190–205.
- . 1998a. *Mind, value and reality*. Cambridge, MA: Harvard University Press.
- . 1998b. *Meaning, knowledge, and reality*. Cambridge, MA: Harvard University Press.
- Pagin, Peter. 2000. Publicness and indeterminacy. In *Knowledge, logic and language. Questions for Quine*, edited by Alex Orenstein and Petr Kofárko. Dordrecht: Kluwer, 163–80.
- Peacocke, Christopher. 1998. The modality of freedom. In *Current issues in philosophy of mind*, edited by Anthony O’Hear. Cambridge: Cambridge University Press, 349–76.
- Quine, W. V. O. 1960. *Word and object*. Cambridge, MA: MIT Press.
- . 1969. *Ontological relativity and other essays*. New York: Columbia University Press.
- . 1990. *The pursuit of truth*. Cambridge, MA: Harvard University Press.
- Smith, Barry C. 1992. Understanding language. *Proceedings of the Aristotelian Society* 102:109–41.
- . 1998. On knowing one’s own language. In *Knowing our own minds*, edited by Crispin Wright, Barry Smith and Cynthia Macdonald. Oxford: Oxford University Press, 391–428.
- Strawson, Galen. 1996. *Mental reality*. Cambridge, MA: MIT Press.
- Wittgenstein, Ludwig. 1953. *Philosophical investigations*. Oxford: Blackwell.
- Wright, Crispin. 1989. Wittgenstein’s rule-following considerations and the central project of theoretical linguistics. Reprinted in Wright (2001), 170–214.
- . 2001. *Rails to infinity*. Cambridge, MA: Harvard University Press.