

Symposium on Belief and Meaning\*  
(Philosophy and Phenomenological Research vol. LVIII, no.3, September 1998)

Akeel Bilgrami

Summary of Belief and Meaning

To put it as broadly as possible, my book tried to provide a theory of the content of our beliefs (and meanings of our utterances)<sup>i</sup> that meets two constraints.

*First*, a constraint that emerged out of Frege's consideration of a certain puzzle about identity, later formulated explicitly as a puzzle about belief by Kripke.<sup>ii</sup> This requires that beliefs attributed to an agent should make her rational by her own lights --in particular, if she is merely ignorant or misinformed (about the identity of certain planets or cities, say), the beliefs ascribed to her should still represent a completely consistent state of affairs, i.e., the world as she conceives of it.

*Second*, the beliefs ascribed to her should not be completely at odds with her external circumstances, or more generally her psychology should not be characterized in a way that allows for radical scepticism about the external world. This second constraint generates the doctrine which I, following others, called 'externalism', a denial of 'methodological solipsism'.<sup>iii</sup>

On the face of it, these two constraints pull in opposite directions. The second constraint requires that the content of our beliefs are in substantial ways determined by circumstances external to us, for that alone would disallow radical scepticism about the external world. But the first constraint demands that the contents of our beliefs reflect the world *as we conceive of it*, that they are determined solely by our conceptions of the external world and not, to put it awkwardly, the external world itself. It is thus no obviously easy task to satisfy both constraints.

Many philosophers have not thought it possible to satisfy both constraints, and thought instead that we need two notions of the contents of beliefs, one that satisfies the first constraint, the other which satisfies the second. They have called these two notions respectively 'narrow' and 'wide' content. Actually this makes it sound as if they were explicitly from the beginning set upon satisfying both constraints. That is misleading. Rather the tendency among philosophers was this. They simply tied content to external circumstances in accord with certain

widely held accounts of reference. (These are a family of doctrines which go under various names such as ‘direct’ reference as in Kaplan, ‘causal’ accounts of reference as in Kripke and Putnam, or more social versions of these earlier accounts as in Burge and Kripke’s Wittgenstein. I will discuss all these doctrines under the label ‘orthodox externalism’.)<sup>iv</sup> They then, as an afterthought, as a result of certain puzzles like Frege’s and Kripke’s puzzles which arose for their orthodox externalist accounts of content, began to see the point of the first constraint, and many of them still downplayed its central importance and the genuine difficulties of meeting it. And when they did pay explicit and detailed attention to it, they offered the bifurcated theories of ‘wide’ content (a notion of content based on these orthodox doctrines) and ‘narrow’ content (a notion of content wheeled in to solve the puzzles and therefore meet the first constraint).<sup>v</sup>

My strategy in the book was not a variant of this. Instead, I brought the first, Fregean constraint to centre stage right at the outset, and then asked what can be said of externalism, if this constraint is met. The idea was not to retreat to *a second notion of content* (‘narrow’) in the face of difficulties arising for a commitment to a very specific orthodox externalist doctrine about content (‘wide’) owing to specific commitments to current doctrines of reference. The idea was to make no commitments to any specific version of externalism from the start, and therefore to see if *a single notion of content* could satisfy both constraints at once. This is what I called a ‘unified’ notion of content.

Two closely related radical conclusions of this strategy emerged when it became clear to me that a) any notion of content that meets the first constraint, can then meet the second constraint only if it abandons orthodox externalism. That is, only if it abandons all standard notions of reference as being relevant to content and meaning altogether. For, as I tried to show at length, it is these notions of reference which were indifferent to and unable to capture an agent’s conception of the world as demanded by the first constraint. And therefore that b) meeting the second constraint now required that we fashion a new externalism which unlike the orthodox one does capture an agent’s conceptions of things.

To fashion a new externalism I imposed a requirement on the way in which external circumstances determine the content of our beliefs and the concepts that compose them.<sup>vi</sup> The details of this requirement and of the new externalism it generates will be impossible to relate here, but an example may convey the very general point. Suppose there are two agents one of whom knows chemistry and the other does not. The requirement is that this difference should make a difference to the way in which a certain external substance determines the concept which they express with their term ‘water’. In other words, though the external substance determines the concepts (and therefore contents) they both have, it determines different concepts if the background of beliefs of the agents are different. Because the external substance is relevant to the determining of the concept one remains an externalist about content (if there was no external

world with which we stood in causal relations, we would have no concepts at all), but because these determinations are restrained by background beliefs of particular agents, these contents capture the external world as the agent conceives of it.

This externalism was not based on the orthodox views of reference, none of which were sensitive to this requirement. And so reference in this standard sense simply dropped out of the study of meaning and content.<sup>vii</sup> This was bound to create a lot of doubt in the minds of a longstanding orthodoxy in the philosophy of language and mind. I raised many of these possible doubts myself and answered them as scrupulously as I could.

The most important of the doubts that I addressed explicitly was this. On the face of it, contents and concepts became very fine-grained and messy on my picture, which they are not on the orthodox externalist picture. This is because different agents were bound to have very different background beliefs and so it would turn out that agents would hardly if ever at all share concepts and contents. Almost everybody would express different concepts with their term ‘water’. (From now on I will say his “ ‘water’-concept” instead of the longer phrase “concept expressed by his term ‘water’ ”.) This would make content and concepts very messy, a distinct disadvantage when contrasted with orthodox views in which concepts are determined by standard theories of reference where everybody shares the same concept if they stand in external causal relations with the same substance, since there is no relevance that background beliefs have to the referential determinations of concepts on the orthodox picture. To tidy up this seemingly messy consequence I pointed to a very obvious but curiously neglected feature of content ascription, which I called the ‘locality’ of content. Locality is a form of contextuality. Contents I claimed were ascribed in particular contexts only. Let us work with examples again to explain this. Take our two agents again, one chemically knowledgeable the other not. Let us say that the former agent is drinking a certain substance from his kitchen tap because he is thirsty after a game of tennis. In this context his background of chemical beliefs<sup>viii</sup> are irrelevant to his ‘water’-concept. So, in context, he may well share his ‘water’-concept with the latter agent. Indeed in the context of explaining thirst-quenching actions of this kind, all or most agents may share their ‘water’-concept for we may abstract from *differences* in the background beliefs since they would all be irrelevant to this action. So, in context, things are as orderly and tidy for my view of content as they are on any other view of content. But context or locality is vital to order.

However what is equally important is that order is achieved without the failure to meet the first Fregean constraint on content. It is achieved within the capture of an agent’s conceptions of things, and not at the cost of it. This is a distinct advantage over the orthodox versions of externalism. For in a context where a chemically knowledgeable agent’s background of chemical beliefs *is* relevant to his actions, such as those say in a chemical laboratory, we

should want his content to *not* be the same as the chemically ignorant agent. This is for exactly the same reason why we want Kripke's Pierre to have different concepts of London when he is in Paris and when he comes to London. It is to capture his conceptions of things. Thus once the crucial idea of locality is in place, my view meets the Fregean constraint *without* making contents messy.<sup>ix</sup> The idea of such locality, such radical contextuality, however, seems to many to go against the very idea of a *theory* of content and meaning. Theories, they would think, ought to say more general things about the phenomena they are studying, things about what concepts and contents a person has, independent of context. And my response to this is simply to say that the locality thesis is so obviously true of our linguistic practice and representational experience that if we have embarked on theoretical enterprises that were defined upon denying the obvious, then we should be glad to relinquish those theoretical projects. In this respect, my claims had the flavour of Wittgensteinian deflation, though I will not claim that the vast mass of avowed followers of Wittgenstein have thought of the deflation along exactly or even approximately these lines.

There were four other doubts I addressed.

Three of them were closely related. Like the previous doubt, all these three took the form of saying that the orthodox externalism had an advantage over mine. Each doubt said that there was something about content and meaning which my view could not capture but orthodox externalism did. The first was that it could not capture an essentially social aspect of language and meaning. The second was that it failed to capture an essentially normative aspect of word-meaning. And third that it failed to capture a distinction that is very intuitive, a distinction between theory-change and meaning-change. I will discuss all of these doubts, elaborating in most detail on the second of them since it is central to the first and the third.

1) The sense in which meanings are not social, I argued, is this. An individual need not be attributed the concept of his fellows or of the experts among his fellows, if he did not have the expert knowledge. So, to take a well-known example, if in our English-speaking society there was an individual who did not know that arthritis was a disease of the joints only, then his 'arthritis'-concept need not be thought of as Burge and others thought of it. There could be individual variance from the experts, as far as concepts were concerned, if the locality demanded it to best specify the agent's conception of a disease. I argued that there might be good epistemological motivations to think of meaning as public, but that did not mean that meaning had to be social in Burge's sense. Individual meanings even when at variance with the meanings of others in the society were perfectly public, since others could and did understand these idiosyncratic meanings. Thus I argued that the private/public distinction simply did not coincide

with the individual/social distinction. The idea that individuals' meanings or concepts were not constituted by the socially constituted reference or extension of arthritis in Burge's sense, may give rise to the anxiety that there would be no norms of meaning. After all if someone in our community said "I have arthritis in my thigh" and could be counted as saying something true given what he means by arthritis, rather than as saying something false and meaning what we (and our expert doctors) mean by arthritis, then there is only difference in meaning, no mistakes in meaning. And the possibility of mistakes in meaning is what makes for meaning being normative. So, following Davidson's denial of the relevance of conventions to word-meaning,<sup>x</sup> I then also denied the relevance of norms to word-meaning and used this to dismiss not only this anxiety about norms but also the entire problem of the sceptical paradox about *meaning* that Kripke attributes to Wittgenstein. This needs some elaboration since it is bound to raise controversy.

2) Wittgenstein claimed that intentions were normative in the sense that an intention divided the world of actions prompted by those intentions into those which were in accord with the intention, and those which were not. "Accord" and "failure to accord" here were meant to capture the normative element. Kripke applies this point as a defining feature of the meanings of words like "plus", "table", etc. It is used as a premise by him in order to rule out certain obvious answers to his sceptical paradox about meaning, and then to propose an account of what allows for the normative element in word-meaning which is a version of the socially or communally constituted reference that I earlier described as one of the orthodox externalist accounts. I rejected the account by rejecting the premise itself.

According to this premise, if I use words with intentions then there must be mistakes of meaning possible since some uses could fail to be in accord with the intentions with which I use words. Thus it is said that if I say "That's a fly" pointing to a black spot, then I would have failed to use the word in accord with my meaning-intention for that word, which is the intention to apply the word 'fly' to flies.<sup>xi</sup> But this is a confusion. The intention that is directed to meaning when I used the word "fly" in that utterance is not the intention to apply it to a fly and nothing that is not a fly. The latter is of course an intention of mine but it is not the intention that is directed at the meaning.<sup>xii</sup> The intention directed to meaning when I use the word 'fly' is the intention that I say something which has the meaning \*fly\* and nothing but that meaning. So for instance if one thought that meaning was given in terms of truth-conditions, then when I say "That is a fly", I say it with the intention of saying something which has the truth conditions that that is a fly. This intention is very different from the intention to apply the word 'fly' to a fly, which of course is also an intention I might have in making the utterance --but which is not directed to the question of the *meaning* of 'fly'. The norm attaching to this last intention is indeed violated if I say "That is a fly" pointing to a black spot. But the norm attaching to the

intention directed at meaning, i.e., the semantic intention, is not violated in these circumstances. For even if I say “that is a fly” pointing to a black spot, my action would still accord with the intention of saying something with the meaning or truth condition that that is a fly. It’s just that the truth-condition has not obtained.

The question arises whether any intentional action in which I uttered the word “fly” could possibly violate the intention relevant to meaning with which I uttered it. I denied that this could happen. This was a radical claim. In order to establish this I considered slips of the tongue in which I might say, for instance, “I am going towndown”. I argued that if this was said with the intention of saying something with the meaning or truth conditions that I am going downtown, then that is what the words meant. Thus ‘towndown’ on my lips then meant downtown. Moreover that was its literal meaning, not its pragmatical interpretation. It was not like a case of metaphor like “Man is a wolf” where the literal and intended meaning or truth condition is that man is a wolf (a condition that does not obtain), and it is exploited so as to yield a pragmatical interpretation which is something else, like perhaps that human beings are competitive. In the example above the *literal* meaning is “I am going downtown.” This, thus, is an example of how there is no mistake made at all except that of misspeaking, of *uttering sounds* that one did not intend to utter. No mistake about meaning. The point of the example is to bring out that one cannot say things that fail to accord with those intentions of ours which are relevant to meanings. So normativity of the sort that is a premise in Kripke’s generating a sceptical paradox about meaning is not something that has any purchase.

I did not deny that in order to mean various things idiosyncratically as in the example just discussed, that is in order to successfully have meaning-intentions that are idiosyncratic, there are bound to be a background of regularities in one’s use of words. One cannot just intend to mean anything and expect to be understood. But the idea that there are regularities in the use of words which make intending meaningful utterances possible does not amount to normativity. Norms attach to the use of specific words. Norms say *this or that* word should be used in the following way and not be used in any other way. But the fact of regularities yield no such norms. We may use this or that word quite as we please and be understood quite well, so long as there are regularities in the background. I also did concede that there might be more utilitarian norms such as “Use this or that word in this or that way and no other if you want to be understood *easily and without causing strain* “, but these are norms that are not constitutive of meaning in the way that Kripke’s premise demanded.<sup>xiii</sup>

I also showed how my denial of norms of word-meaning ought not to be confused with the denial of norms regarding logical and rational relations between sentences and relations between beliefs. One may deny norms in word-meaning while insisting that intentionality was constituted by rational norms. It was Kripke’s use of the normative element as being relevant to

particular concepts, to the meanings of words like ‘plus’, ‘table’ etc. that I was denying, not his general distinction between normative aspects of intentionality which are to be contrasted with mere dispositions. In a curious sense, then, my diagnosis of Kripke’s mistake is that he changes his own good subject of normativity in the very idea of intentional states to a subject having to do with the meanings of terms. Because, despite Kripke, there are no mistakes in meaning, however, it does not mean that he is not right that intentionality is a thoroughly normative phenomenon. What this means is that one must be much clearer than he has been about the relation between his two subjects, intentionality and word-meaning.

3) As for the third doubt about the distinction between theory- and meaning-change, it really extended this anxiety about finding something ‘wrong’ (normative) versus finding something merely different (non-normative, descriptive) to the following question: if meanings of theoretical terms were not given by reference, a succeeding theory would not count as a change of *theory* from some previous theory, but merely a decision to adopt new meanings. If there is to be theory-change instead of mere meaning-change, then there would have to be some common subject matter between the two theories, and that could only lie in the common references of some of their theoretical terms. I cannot retell my whole response here but to simply assert that I argued that in any particular case in which we felt that it was a case of theory-change and not meaning-change, that could be captured quite adequately in terms of shared beliefs between the previous theory and the new theory (these constituting the common subject matter) and it required nothing like causal and direct accounts of reference of the terms of the theory, as many philosophers thought.<sup>xiv</sup> If someone claimed that there is theory-change and not meaning-change, even if there are no shared beliefs between the two theories, I argued that this was dogmatic, and only a philosopher already in the constricting grip of an orthodox referentialist view would even be tempted by such a claim. It was not as if the claim provided an argument for the orthodox referentialist view, it was a prior commitment to the orthodox view that prompted this dubious claim.

Finally, I addressed a fourth doubt about my specific formulation of the second constraint regarding externalism. My formulation of the constraint was in epistemological terms. In this it followed a formulation of Putnam’s early article “The Meaning of Meaning” where he defines externalism simply as the denial of the claim known as methodological solipsism, a denial of the claim that our psychologies "do not presuppose the existence of any thing external to the agent who possesses such states." (p.136) However the fact is that Putnam and most others had all sorts of *other* motivations also for wanting externalism as a constraint on content, such as those just mentioned in the previous many paragraphs. To repeat: meaning and the mental must be ‘social’ and ‘normative’ in very special senses of those terms, and meaning must be such that we

can distinguish between meaning- and theory-change; and only an externalism of a very specific (orthodox) sort will fulfil these motivations. As I have said, I regarded these further motivations as either bad ones, or when good (if properly understood), fulfillable by a quite different externalism, one which *unlike* the orthodox one was able to *also* meet the first Fregean constraint. That is why I repudiated all these motivations and returned to the initial motivation for externalism in Putnam, the denial of methodological solipsism or (what is widely and somewhat abusively these days known as) Cartesianism. So the last of the doubts I addressed against my account is one which wonders why epistemology should be brought into the study of meaning and belief. The other motivations having to do with norm, society, and meaning-change/theory-change were at least all motivations having directly to do with the nature of meaning and content. Epistemological considerations about scepticism have no such direct bearing and seem to some to be unnecessary things to bring in as constraints on content and meaning. In response, I argued in detail the importance of epistemology to the ideal of the public nature of meaning and thought, and in turn defended the ideal of publicness of meaning and thought as being central to a credible general epistemology. I showed how those who had discussed internalism versus externalism in the study of content had willy nilly taken on an epistemological issue. And I showed how the standard disputation on this subject (such as the one invoking brains-in-vats) was question-begging on both sides, and tried to show how we might find ways of avoiding begging the question.

But at a more general level my first response to this doubt about the role of epistemology as being relevant to the study of meaning and intentionality was to say that it betrays a lamentable lack of ambition. Surely the investigation of language and mind (of representation in general) is a means to advancing our understanding of the world in metaphysical and epistemological terms, and that is certainly how it was first motivated by those who pioneered these subjects. Now that the study of language and mind has developed into a self-standing subject, and it certainly has by now, it would seem a shame if the initial more general and deep philosophical source of concerns behind them got lost. Thus by formulating my second constraint in the study of belief and meaning in epistemological terms, I was raising the stakes that had been lowered in the less ambitious pursuits of recent years (pursuits which of course were highly fruitful, instructive and worthy within their own narrower concerns). I should not give any impression that this was an innovation on my part. Donald Davidson in some of his later papers (following the mature Wittgenstein) had raised these stakes as had others who had been influenced by him such as McDowell and Evans, all of whom had explicitly re-introduced epistemological considerations of anti-Cartesianism in the study of language and mind.<sup>xv</sup>

So a good bit of my book addressed itself first to the importance of these epistemological considerations in Davidson and in McDowell and Evans, and then criticized them for not having

satisfactorily responded to these considerations. Given these shared motivations, naturally my disagreements with them were much more subtle and more internal to a shared philosophical framework for asking these questions, than my disagreements with what I am calling orthodox externalism. Davidson, I argued, simply had not paid enough attention to the first constraint and not allowed himself the resources to meet it by formulating the role of truth in the study of meaning along the lines he did. My suspicions about this inadequacy in his work has been confirmed in recent writings of Davidson where he very explicitly says things which deprive him of all possible resources to meet the first constraint.<sup>xvi</sup> In the case of McDowell and Evans my criticism was that the externalism by which they met the second constraint --their famous (infamous) Russellian thought thesis, whereby there could be no singular thought in the case of such intentional episodes as hallucinations-- was both counter-intuitive and unnecessary, but worse that externalist content defined along these lines did not quite meet the demands of the first constraint in a way that genuinely captured an agent's conception of the world.

The book thus takes on a lot of comers in the contemporary field on subjects that are very widely worked-over and hotly contested. It argues for a unified notion of content against a widespread tendency to split content into two. It argues that certain standard theories of reference are altogether irrelevant to meaning and content. It argues that meaning, concepts and contents are radically local and contextual and therefore certain familiar projects going under the name 'theory of meaning' trump up a subject that is not really there. It argues that many of the standard ways of thinking of language as social, meaning as normative, and the relation between theory- and meaning-change are all misguided.... It's not surprising then that this symposium on the book has registered some strong and interesting counter-argument. I turn now to my four commentators and critics, to all of whom, I am most grateful for having taken the time and effort to write these comments.

## Replies

**Michael Williams's** challenging and interesting comments bear on an absolutely central issue I raised, the issue of the first constraint. The puzzles I considered from Kripke and Frege force upon us the need for such a constraint, the need to capture an agent's conceptions of the world (of cities, planets...). Williams says that it is the strength of Davidson's position that he is unconcerned by these things because he thinks that the very idea of conceptions is a holdover from a way of doing philosophy that was pre-Quinean where philosophers still took seriously the

idea of conceptions as doing serious explanatory work. The very question: does Pierre have a different conception when he expresses things with the term “London” from those he expresses with his term “Londres”? is for Davidson a bad question. Conceptions, senses, all such notions smack of analyticity and are outmoded.

So the question arises: How does Davidson who thinks that meaning is to be understood in terms of truth-conditions, respond to Kripke’s puzzle? According to Williams it is not by writing down two different truth-theoretic clauses for the relevant fragment of Pierre’s language: “‘London’ refers to London” and “‘Londres’ refers to Londres”. This would be to fall prey to the idea of (two different) conceptions of the city. And all talk of conceptions is bad. But how is the puzzle to be resolved for Davidson, if not this way? Williams response is: “Davidson’s holistic methodology leaves the boundary between meaning and belief fluid; there are always ways of trading off sameness and difference of reference against differences in belief. So not only does everyday talk of concepts fail to yield clear identity conditions, the interdependence of meaning and belief, by excluding the analytic-synthetic distinction, shows why such conditions should not be expected.”

So it would seem the puzzle is resolved for Davidson by saying that there is no reason to say in the theory of *meaning* for Pierre that “London” and “Londres” are two different concepts (or meanings) by writing down in the clauses of the truth-theories for these two terms two different right-hand sides, as I did. That would be to foolishly take talk of conceptions seriously. Rather Davidson simply points out that in the beliefs attributed to Pierre we just don’t attribute the belief that London is Londres. And that would make clear that there is no puzzle. But this cannot possibly be right. The theory of meaning is the only resource we have for specifying what goes into the contents of a person’s beliefs. So if Pierre were to go on to wonder if London is Londres (something that he well might if he does not actually believe that London is Londres) there would be no resources in the theory of meaning to show how he could possibly wonder this. For we must ask what is he wondering this about? What do “London” and Londres’ in that wondering signify? If we turn to the meaning-theory for him (as Davidson and Williams think of it) to answer this, we find that his wondering would be as idle as my wondering whether bachelors are unmarried. But, as is well known, it may be no trivial or idle wondering on the part of someone who does not know the identity. The point is simply this. Even for Williams who says the puzzle is easily resolved by looking to Pierre’s beliefs and finding he lacks the crucial belief about an identify, we still have to have some idea about what goes into the specification of Pierre’s beliefs. Whatever that is will have to distinguish between “Londres” and “London” if such non-trivial wondering is to be possible. And so we will need some way of counting these as *two* concepts. (If one does not like the word ‘concepts’, one should chose another word.) And it was my claim that the two different right-hand clauses of the theory of

meaning keep track of the fact that there are two different concepts here which account for why Pierre can non-trivially wonder this. Davidson and Williams cannot, then, simply pretend that serious issues forced by Kripke and Frege's puzzle can be swept under the carpet because they don't like the rhetoric of "concepts". If pointing to Pierre's beliefs won't resolve the puzzle, then they owe us an answer as to what will. I submit that nothing will unless it is established that Pierre's two terms mean different things (express two different concepts).

Williams says that since Quine's attack on analyticity we should know better than to talk of conceptions, implying that my talk of conceptions implies some sort of commitment to analyticity. My insistence on locality was explicitly introduced by me partly to stave off anything like analyticity from my talk of an agent's conceptions of things. Take the example of 'water' again. I argued that to the extent that one thought that there must be some belief (or cluster of beliefs) that *had to be* selected in *every* locality in which someone had a 'water'-thought in order to characterize the 'water'-concept in that locality, for example the belief that it was H<sub>2</sub>O, or that it was transparent, or .... then to that extent one believed in analyticity. But I precisely *denied* that any particular belief (or cluster of beliefs) had to be selected in each and all localities in characterizing a concept. Thus the notion of sense or conceptions I invoked to solve the puzzles could not be convicted of hankering after analyticity.

**Gary Ebbs's** commentary contains a miscellany of points, some of which centrally address my views.

Before I respond to them, I should say at the very outset that there is an extraordinarily misleading and inaccurate use of a quotation from Davidson in Ebbs's comment which altogether misrepresents the point Davidson is making. The quotation is made in the context of discussing my view that the simple statements of truth conditions derived from a truth theory for a language do not say all that has to be said about what goes into the interpretation of a speaker. Ebbs cites a remark from Davidson about me which has nothing to do with this claim of mine, in order to show that Davidson disagrees with me about this claim. Davidson's remark has to do with something I had written about how his principle of charity is more generous than what "minimal" rationality requires in interpretation. It has nothing *whatever* to do with my claim that Ebbs's is discussing. And though Davidson might disagree with me about my claim (something that Ebbs's might think again about since Davidson also claims that a theory of truth has to be linked with a statement of a speaker's beliefs and desires before it can serve in interpretation), it is quite deliberately misleading to say that he does so in this remark.

I turn now to the various points in his comments.

The first point is that I do not distinguish between concepts and conceptions. This is a familiar distinction. Roughly speaking those sympathetic to orthodox externalism have often

said that denotation or extension yields the concept but how we conceive of the object denoted or the objects that fall within the extension, yield the conception.

To repeat, according to orthodox externalism denotations and extensions are central to explicating our concepts of cities, planets, substances, diseases, etc. I, following a legion, had made much of the fact that this is what on occasion gave rise to the sorts of inconsistencies that Kripke addresses in his puzzle about belief. I had also raised a similar difficulty not just about individual concepts<sup>xvii</sup> or proper names, but about our concepts of substances, diseases (“water”, “arthritis”) etc. I had said that if a (merely) ignorant English speaker said “Water is not H<sub>2</sub>O” or “I have arthritis in my thigh”, then, *on the orthodox externalist view*, this led to attributing inconsistencies similar to the one Kripke raises. It thus violates my first, “Fregean”, constraint. If extension characterizes these concepts, then there is some reason to think that the speakers here are saying and (given the principle of disquotation) believing inconsistent things. Ebbs spells out my point with an example from Burge that I discuss. He spells it out in the belief-ascriptions --numbered 1 to 3 in his comment-- to a protagonist called Bert. He says that if I had made the distinction between Bert’s concepts and Bert’s conceptions I would have seen that no inconsistency is implied. In his words: “Once we distinguish between concepts and conceptions, we can accept (1) without committing ourselves to (2) or (3).” Let’s take (2). Ebbs says what he says presumably because he thinks that I think that “**rheumatoid ailment that afflicts the joints only**” in (2) is Bert’s conception substituting illegitimately for “**arthritis**” in (1) which is his concept. I had in fact discussed the question of whether this substitution is legitimate or not in my book and by the end of that discussion considered the very position that Ebbs embraces, i.e., that there is nothing we can say about concepts.

But before getting to that, I should stress that Ebbs just gets my view wrong about what is involved in my insistence that accepting (1) requires that we accept (2). That insistence simply does not come from anything to do with my penchant for conceptions and so no amount of distinguishing between concepts and conceptions is going to help with dissuading me from it. In fact it is precisely my point that “**rheumatoid ailment that afflicts the joints only**” cannot be Bert’s conception since he is ignorant of that bit of medical knowledge. It is my point (though see the caution and its implications mentioned in the footnote attached to the end of this sentence) that the orthodox externalist must accept (2) if he accepts (1), not because he is aspiring to get Bert’s conceptions but because he thinks the concept is characterized in terms of the extension and (2) captures the extension.<sup>xviii</sup> But let’s put this basic mistake of interpretation aside. Perhaps Ebbs meant to say that if you foisted on the orthodox externalist a view that terms expressing the criterion by which the extension of a term is fixed must be substitutable for the term itself in content-ascribing contexts, then you had somehow converted concepts (which are to be thought of in terms of extension) to conceptions. If that is what he meant to say, it is

certainly not my view of (Bert's) conceptions since for me there can be no conceptions of which their possessor is unaware in some broad sense of 'aware', and Bert is in that sense unaware of what Burge and Ebbs insist is the extension of his term; and all this is why I had in fact spent some time discussing what exactly it was that I was foisting on orthodox externalism in making that substitution if it was not my view of conceptions, and why I was right to do so.

Even if Ebbs misinterprets me in order to work up to the position he upholds, what he upholds is that Bert's concept cannot be something that we can say *anything* about, so there should be *no* substitution of anything, whether of something that captures his conceptions *or something that captures the extension*.. But this means that a concept is not even characterized in terms of extension, for the idea of the extension of water does allow us to say something about it, i.e., that water is any substance that has the chemical composition H<sub>2</sub>O. (It may be hard to see what then is left of the externalism that I was criticizing, but let that go.) On this view nothing can characterize the concept. That is the view which in my book I had said made concepts a mystery.<sup>xix</sup> Now, I confess that I am silenced by someone (Ebbs) whose response to this is: "I love a mystery!" That is certainly a conversation-stopper. But that is precisely why I had invited the reader of my book to explore the possibility that we might try and do better and consider other views of concepts than the one that insouciantly says that we cannot say anything about concepts. Besides I had also discussed at length (pp.38-41) the distinction Ebbs is making between concepts and conceptions and had attributed it to Burge, and during the discussion had given reasons for rejecting it, reasons quite other than that in a certain version of the distinction (the one that Ebbs's embraces) concepts were defined by this insouciance. I had also given reasons as to why in *any* version of the distinction, so long as the notion of concepts was defined in terms of orthodox externalism, that notion was not well motivated in the first place and, as I described in my summary above, I considered all these motivations in detail and showed them to be bad ones. Ebbs does not address any of this, so I will leave the matter there.

As a general point, I should also make clear that Ebbs is quite wrong to think that my objection to orthodox externalist views of concepts and content come from *intuitions* I have about explanation. They come from an entirely theoretical aspiration to meet what I called Frege's constraint in my summary, constraint (1). This is not an intuition of mine. It is a theoretical constraint that many have placed ever since Frege and what I pointed out was that the constraint's demand was a demand about explanation as well. This too is something that is quite widely demanded (see footnote 7 above). Many of those who have made this demand have unlike me *not* gone on to reject the orthodox externalists altogether and have been happy to introduce a second notion of content to meet the constraint and capture the psychologically explanatory elements that orthodox externalism fails to capture. But this is not Ebbs's position apparently. He is not moved in the first place by the theoretical demand, which he seems to see

as misplaced “intuitions” issuing from a misguided desire to say something rather than nothing about concepts.

Ebbs then raises some questions that are intended to be addressed to me but are questions that can be raised for many views that are, very generally speaking, like mine, and to which the answers, whether effective or not, are so well known that there is not much point responding to them in detail here. One is a question regarding my qualification on external determination of concepts that it respect the other beliefs of an agent, and not do it as traditional causal theorists do, purely in terms of causal chains or causal co-variances with items in the world. Ebbs says that this presupposes that the interpreter *already* knows what the other beliefs are, with the suggestion that this is to assume too much. This is a perfectly general problem for radical interpretation or any method that is both holistic *and* has to start somewhere in particular. And the answer is well known: as an interpreter you assume by way of hypotheses what a person’s range of beliefs are (you may if you like assume that they are the ones you believe), and if you are wrong in this, then that will show up in failing to save the phenomena (the data) in particular cases, and so you go back and change some of the initial hypotheses and see if things go better, and you proceed in this way, until a fuller theory develops with more systematically linked and confirmed hypotheses. Thus in our example in particular, you may assume that someone has the chemical and all the other beliefs you have regarding water, and then if you find that particular attributions in the presence of what you consider water do not explain the behaviour as you would have expected, you go back and withhold some of the beliefs you had attributed at the outset. Now a new and revised set of background beliefs are posited and a different ‘water’-concept is determined because of that, which may account for the overall behaviour better. This is an application of the standard holistic method not just about hypotheses about meaning but about anything. Expositions of this drill can be found in much greater detail in many introductory books in the Philosophy of Science; Quine explicitly mentions it too and so does Davidson. Ebbs may not find these accounts very convincing. But that would not be directed to my account in particular. My account simply follows others in assuming that the idea of “already knowing” what someone’s beliefs are as you solve for a particular externally determined concept or “indexical” belief is a result of some such method.

He also says I do not give much attention to the method of how compositionality is established for my view as opposed to a standard truth-theoretic view, where it is clear how it is. My view is roughly a combination of a description-theoretic view (see footnote viii) and a conceptual-role view, with the qualifications of locality and various other qualifications I discuss in the book which try and address the standard objections to description theories that have nothing to do with compositionality. Nobody has ever said that a description theory cannot keep

faith with compositionality (nor that a conceptual role theory cannot). I did not therefore pay more attention to the issue than in the passage from me that Ebb cites.

He concludes with some questions about my locality thesis. I had said that in a particular context or locality, not all of a subject's, say, 'water'-beliefs may be relevant in assigning a 'water'-concept to him. For instance a chemical belief may not be relevant to the 'water'-concept in some content that explains someone's drinking water to quench thirst. So in such a context (say where someone is in a kitchen after a game of tennis and we want to explain a particular action of his) we may just stick with some other beliefs such as that it comes out of taps and that it quenches thirst. So let us assume that there is an aggregate of 'water'-beliefs that a person may have which includes his chemical beliefs and these other beliefs, and we select only the latter in this context. Ebbs says 'select' is misleading. How do we know that the belief formulated in the that-clause "believes that water comes out of kitchen taps" (which is one of the many beliefs which go into fixing what the 'water'-concept is at the aggregate level) is the same belief as the one formulated with exactly the same words in the local context (where it goes into fixing what the local 'water'-concept is to explain the action in the kitchen). The criticism seems to be something like this: we cannot tell for sure that the beliefs are the same just because the that-clause formulations read the same; and Ebbs thinks that since --on my view as opposed to Burge's view which appeals to everyday ascription-- the term "water" may express more than one conception, more than one 'water'-concept, there is a significant gap or distinction between what a term expresses and the term. He seems to think that this is a distinction between use and mention, but that shows that he has an imperfect grasp of the use-mention distinction. There certainly is a distinction here and of course one may raise a question about how, once we depart from every day belief-ascription, we can tell which belief some that-clause ascription expresses. But why should this be an insuperable problem? Here is how you tell. Take Burge's Bert. I start by assuming that when he says "I have arthritis in my ankle", he means what knowledgeable people like me and the experts mean by "arthritis", and so by the principle of disquotation I attribute to him the belief that he has arthritis in his ankle. I then observe that he says "The arthritis has spread to my thigh." So I begin to wonder whether he does not mean something else by his term, what in his mind are a wider class of ailments. I may then decide to revise my initial attribution to him. This will of course mean that we must find some notation to distinguish the two concepts (his concept is "tharthritis" and not "arthritis", I might say, rather than merely saying what I have been saying so far, which is that he has a different 'arthritis'-concept than mine.) Neologizing is necessary to represent beliefs, even if we do not do it explicitly but only implicitly. That is to say that something like it is taken for granted. This point is obvious because without it we would not have any clue as to how to solve basic puzzles such as the one Kripke raises about Paderewski. Anyone who thinks that we do not need to assume

neologisms, such as, say, unpronounced subscripts (Paderewski<sub>1</sub> and Paderewski<sub>2</sub>) is never going to solve the puzzle. (The ‘unpronounced’ there indicates the implicitness of the neologism.) Since the puzzle is a puzzle about belief and since it is apparently not insuperable once we introduce the neologisms, that is just proof that one can depart from everyday socially constituted belief-ascription (i.e., we can make ascriptions of beliefs invoking neologisms) and perfectly well answer the question of how to correlate that-clause representation of beliefs with the beliefs they represent. The fact that we can solve these puzzles demonstrates that the contents we ascribe here are publicly accessible even as they clearly depart from everyday ascriptions. And that is exactly right. Departure from the social need not amount to attributing private contents. In my book I was at pains to say that the two distinctions, Individual/Social and Private/Public, are not coincident. Hence if individually fine-grained contents are publicly accessible, they are communicable; so it cannot be the case that we cannot tell that the beliefs in the aggregate are or are not the same as the ones in the locality. And if they are the same then “select” is a perfectly good word for what I do when I take some but not others from the aggregate to the locality.

There are of course other much deeper problems about my locality thesis than the one raised by Ebbs about correlating that-clause representations with beliefs, problems about what the criteria are by which I select some beliefs rather than others (i.e., by what lights I find some relevant in a locality and not others), and also problems about how to assume that some beliefs may be extricable from others in the aggregate to which they might be inferentially connected. These questions imply hard questions about analyticity, the exact nature of holism, and so on, which I have tried to address in the book and other writings.

**Derk Pereboom** raises the first of these deeper questions sharply in his commentary. This is question 3) in his comments. He asks how one knows what to select in a particular locality, how does one know what is relevant to a particular explanation which defines the particular contextual locality. He says that we must give a general formulation as to what will count as relevant rather than depend on particular intuitions for particular explanations. He clearly does not think it sufficient to say that any content which explains behaviour must meet the first Fregean constraint I have laid down. And it is true that that by itself does not explicate what is meant by relevance. Let me then explore a bit this worry about how to think of what is meant by the relevance of some beliefs and not others to given explanations.

On the face of it there cannot be any serious problem about finding out which beliefs are relevant to explanation, at any rate no more than it is a problem to say what it is that is to be explained. If we know what it is we want explained (say, a person’s act of drinking water after a tennis game) then we know when some information is relevant (that he thinks water is available

in the tap in front of him, that he thinks it has the capacity to quench thirst) and when some might be irrelevant or redundant (that he thinks its chemical composition is H<sub>2</sub>O, that he thinks water might be found on Mars) in explaining it .

I have made the obvious initial move. I have appealed to the explanandum in order to answer Pereboom's question as to what intentional states are relevant in an explanation. However the worry might be that I have purchased a quite uncontroversial account of relevance here only by dogmatically helping myself to something that itself needs to be justified. I have helped myself to the idea that it is transparently obvious what the locality-defining explanandum is. So it might be said that the explanandum is up for grabs too. And if that is so, then we need to know the criteria by which we decide which locality we are in for which we are selecting contents. Fair enough. But all one needs to say in response to that, is this: any judgement as to what the explanandum is is itself an hypothesis, and therefore is itself revisable. So we are not helping ourselves to anything dogmatically. Like all theoretical judgements we may revise these judgements about what the explananda are, if we find we were not getting an overall coherent account of the agent in question. This is a standard and general methodological canon, nothing specific to theories of a person's meanings and contents. but applicable in the study of anything.<sup>xx</sup> The important underlying point is that, relative to a (revisable) decision about what the explanandum is, we know what to select and what to leave out in the concepts that compose the contents in the explanans. Thus there is no serious problem about relevance of the sort Pereboom raises.

Remember the deeper point that locality was introduced to address was that if different agents may have different (say) 'water'-concepts then content would be a very messy notion, since there is a danger that agents may never share any concepts and contents. This is the danger that there might be no generalizations about them. In that case, we could never say they do the same things for the same reason. But they manifestly do. And it is this manifest fact that needs to be captured and I claim is captured by the locality thesis, because despite differences in peoples' 'water'-concepts at an aggregate level, in localities they may be seen as often doing the same thing (drinking water) for the same reason (quenching thirst). Pereboom's question about relevance, by implication, asks how we come upon these shared explananda and explanans.

As I said, the situation on this matter is no different in intentional explanation than it is in explanation of many non-psychological phenomena. So that is why it seems to me that the problem about relevance that Pereboom raises ought not to be one that is restricted to my claim. It is a widespread problem affecting all explanation and not just explanation in psychology, and there are standard ways of dealing with it. My claim is that in psychological explanation, though it is right to demand that people share beliefs and fall under the same generalizations, it is no more difficult to find this generality of description for explananda and explanans here than it is in

the case of many non-psychological explanations. No more difficult, say, than finding a generality of description whereby a fire in the house and a fire in the dock can be seen as throwing up common explananda and explanans. Just as we abstract away from the chemical beliefs in the case of my protagonists, we can abstract away from the fact that in the one case it was the cloth in the furniture which caught fire and in the other the wood in the boardwalk which caught fire. We do not find abstraction to a common general description in this case giving rise to worries about relevance, so why should we think relevance an insuperable problem in the case of psychological explanation, which involves essentially the same sort of abstraction? There is no special problem here for content that is not there for combustion.

To change the subject, the first part of Pereboom's comment offers a fair diagnosis of the differences between my position and orthodox externalists like Burge. He says that in the usual cases of content attribution where there is no gap in crucial (say, medical or chemical) knowledge between an individual and his fellows in society, Burge's view may suffice but in the more challenging cases --where there are gaps of knowledge-- one may have to move to my position, and the latter are the real tests of what the right account of content is. If this is right, it brings out all the more why I oppose Burge since the challenging cases are raised by Burge himself. But there is something that Pereboom says that seems to me to not get quite the right gloss on his own diagnosis. He says that Burge's views are less appropriate than mine when it comes to psychological explanation but they are more apt than mine when we look to the communicative point behind content attribution. What is meant by the communicative point behind content attribution? Let's say someone were to attribute aloud a belief to X, a chemically ignorant person, "X believes that there is water in that jug" and the hearer of this remark Y is a chemically knowledgeable person. Then the communicative point is that Y would go away with some information about the world, viz., that there is a substance in the jar whose chemical composition was H<sub>2</sub>O. Pereboom thinks that this shows that one point of belief attribution (conveying information about the world) is only fulfilled if we assumed that Burge and others were right in thinking that "water" in that attribution must convey that socially constituted (the chemical) concept to X, despite his chemical ignorance. But that does not seem to me to follow. As I say in the book<sup>xxi</sup> that piece of communication of information about the world could just as easily be analyzed as follows. "Water" in that attribution attributes something other (if you like, something less) than the chemical concept to X (since X lacks chemical knowledge) but Y (who has the chemical knowledge) can bring to what is communicated, *her own knowledge* that the substance is H<sub>2</sub>O, even if she knows that X does not conceive it that way. This last step is what gives the hearer the full information about the world in the communicative exchange. On this analysis, what is attributed to someone in psychological explanation and what is attributed to someone in communication to others is the same content; and yet the end result, i.e., that some

information about the world is acquired by the hearer, is still achieved. No need for Burge's more orthodox version of externalism even for the communicative point of content attribution.

Pereboom's question (2) seems to me to be based on a misunderstanding of how I use the expression objects and facts external to agents "constitute" their contents. If I understand the view he himself favours owing to Hilary Kornblith, it is by no means incompatible with mine, though its emphasis is somewhat different since Kornblith (on Pereboom's exposition which is my only exposure to the view) was not primarily concerned with the nature of belief attribution and its rationale.

**Jeffrey Malpas** in his comment registers much agreement with Belief and Meaning.

It's interesting that his real disagreement would turn out to be with Williams and Ebbs since much of their commentary attempted to defend Davidsonian positions on the relation between truth and content against mine, whereas Malpas seems to think that there is no great distance between Davidson and me. I will leave it to them to sort this out.

He encourages me to be more explicit than I was on the relations between agency and intentionality, saying that my linking of intentional content with the themes of action-explanation and self-knowledge to oppose orthodox externalism still leaves things a bit implicit. I would have thought that the last few pages of the book which connect my views of intentionality and self-knowledge with Strawson's on responsibility were explicit efforts at spelling out these links. And subsequent work of mine has pushed these links further than I did in my book.<sup>xxii</sup>

But that apart, I think that Malpas is wrong to think that making these links explicit would silence those who dismiss such views of intentionality as being too anti-realist (or "instrumentalist", as it is sometimes said). That is to say, why would it silence those (Fodor) who think that rationality has absolutely no constitutive role to play in the study of content, to be told that rational agency and intentionality are essentially linked. Or for that matter, why would those (Searle) who think that talk of its role in action-explanation distracts from the 'intrinsic' nature of intentionality which is due to some obscure combination of Cartesian subjectivity and a parochial version of materialism, withdraw the charge of anti-realism against a view that explicitly links intentionality and agency? I myself believe that intentionality is intrinsic but unlike Fodor and Searle my thinking so does not make me reject its link with agency, but rather stress it. For it is precisely the normative character of that link that makes intentionality intrinsic. Perhaps an overly dispositionalist view of what goes into action-explanation is rightly charged with anti-realism and Fodor and Searle are right to oppose that view. But their mistake is to swing from that opposition to thinking that intrinsicness therefore requires delinking intentionality from rational agency and to embed it in denotation alone (Fodor) or, as I said, embed it in an obscure combination of Cartesianism and biological parochialism (Searle). But I

don't think that either Searle or Fodor would be in the slightest bit impressed by my calling this a mistake. Their idea of intrinsicness is simply not mine, and my way of thinking of intrinsicness would seem to them to fail to produce the robustly naturalist realism they seek.<sup>xxiii</sup>

I think Malpas is similarly over-optimistic about the persuasive power (remember 'persuade' is a success-verb conveying perlocutionary effect) of the views we share when he says that there are transcendental refutations of the epistemological sceptic which I am too modest to claim, but would be in a position to claim if I made the links between agency and thought explicit. Nothing in these links could possibly provide a 'refutation' of that kind. At best these links show that intentionality conceived as I do puts aside a premise that is crucial to the sceptic raising his threat. I was clear that that was what I was doing in insisting on the public nature of thought and language and seeing through its implications. These implications refuse to accept a premise of the sceptic, viz., what Bernard Williams has called the "objective conception of reality", which gets a full recent exposition in Tom Nagel's The View from Nowhere,<sup>xxiv</sup> and which I discuss in some detail. If presenting reasons to put aside the sceptic's premise is a 'transcendental' refutation of the sceptic, I would be glad to have made the less modest claim that Malpas urges upon me. But surely transcendental refutations aspire to show that the sceptic's own premises carry in them the seeds of his conclusion's refutation. As Barry Stroud and others have argued, the standard examples of transcendental arguments against the sceptic do not succeed in doing this but rather replace the sceptic's premise with a less fiercely ('metaphysically') realist one.<sup>xxv</sup> That is all I myself claimed to do in devising my epistemological arguments.

Thus on both points, on the nature of intentionality and the related opposition to scepticism, even if I were to make things more explicit as Malpas proposes I should do, I would find it a little brash to describe my claims against the philosophers I find wrong along the lines Malpas suggests. None of that inhibits me in the slightest in my thought that they are wrong

---

\* Published by Basil Blackwell, 1992.

<sup>i</sup> Since space is scarce, I will not keep referring to beliefs and meanings together in this way each time, but rather talk of one or other. The principle of disquotation should make it clear that

---

the two are very closely linked. That is why for instance philosophers like Kripke take issues of reference and meaning to directly affect beliefs in the form of a ‘puzzle about belief’. And that is why philosophers like Burge take certain facts about linguistic usage and meaning to give rise to questions about “individualism and the *mental*.”. (See references below.) One may of course take an interest in a notion of language and meaning that has not much to do with intentionality (as for instance makers of dictionaries do) but it is not obvious that that notion is of great interest to philosophers, nor that it should be of interest when we are considering the notion of meaning that is subject to the principle of disquotation. So I have no (philosophical) interest in a notion of language or meaning which is independent of the notion of what is believed. If the meanings of our words are relevant to generating puzzles about beliefs and repudiating or embracing individualism about the mental, then we cannot also pretend that those meanings and concepts are irrelevant to explaining behaviour and all the other tasks and constraints that apply to beliefs and the mental or intentional. This last point is absolutely essential in motivating the first constraint that follows immediately in the text.

<sup>ii</sup> Frege, “On Sense and Reference”, Philosophical Writings, eds., P.Geach and M.Black (Blackwell, 1970) and Kripke, “A Puzzle about Belief” in Meaning and Use, ed. A. Margalit (D. Reidl, 1970)

<sup>iii</sup> See Putnam “The Meaning of Meaning” in Mind, Language and Reality (Cambridge, 1975)

<sup>iv</sup> Kaplan ‘Bob and Carol and Ted and Alice’, in Approaches to Natural Language, eds J.Hintikka, J.Moravcsik and P.Suppes (Reidl, 1973); Kripke, ‘Naming and Necessity’ in Semantics for Natural Language, ed., D. Davidson and G. Harman (D. Reidl, 1972); Putnam, *ibid*, Burge “Individualism and the Mental” in Midwest Studies in Philosophy, IV, (Minnesota 1979); Kripke, Wittgenstein on Rules and Private Language. (Blackwell, 1982). These have all inspired innumerable articles taking orthodox externalism for granted.

<sup>v</sup> For a very explicit version of this see Brian Loar’s excellent “Social Content and Psychological Content” in Contents of Thought, eds., R.Grimm and D.Merrill. But he only makes explicit a procedure which is very widespread. Incidentally, I should add that ‘narrow’ content is also (and quite differently) motivated by metaphysical considerations having to do with supervenience of the intentional on the *inner* physical elements of those who possess content. I will ignore those motivations for narrow content since I do not believe that the intentional supervenes on the *inner* physical.

---

<sup>vi</sup> I use the term content to be to belief what meaning is to sentence. I use the term concept to be to content what word-meaning is to sentence-meaning. Thus concepts compose the contents of one's beliefs just as word-meanings compose the meanings of one's sentences. All these four notions are deeply linked and integrated in this way so long as one accepts the close relation between belief and meaning outlined in footnote 1.

<sup>vii</sup> I should add that this repudiation of reference is not at all the same as Davidson's repudiation of reference in his paper "Reality without Reference" (See his Essays on Truth and Interpretation, Oxford, 1983). In that paper Davidson is merely keen to point out that certain permuted cases of indeterminacy of reference make no trouble for truth-theories of meaning since at the level of the truth of sentences the theories may be equivalent. My repudiation of reference is much more radical and is motivated by quite different considerations. In Chapter 4 of the book, I describe in detail the disagreements between Davidson and me.

<sup>viii</sup> The reader will notice that I talk of 'beliefs' here where earlier theorists have often talked of 'descriptions'. So where earlier theorists claimed that there were certain descriptions that were associated with a certain term, for instance, 'water', I talk of the beliefs associated with an agent's 'water'-concepts. I think to do this is just to make explicit what was going on in earlier theorists. So in a sense my position may be described as a version of description theory. However I changed two things that were crucial to the earlier theorists. First I insisted that the descriptions for the same term could be very different in different localities. And second I had no role for orthodox notions of reference. So it was not as if I appealed to descriptions (or beliefs) over and above something like reference as orthodoxly conceived. Mine was a description theory of *meaning* and *concepts*, a theory however which *also was externalist* (in a more unorthodox sense) and which captured an agent's conceptions of *the world*.

<sup>ix</sup> It would be obtuse then to take my of account of content and meaning, which aspires to capture the fineness of grain in an agent's conceptions of things, to have failed to address the possibility that that leaves content and meaning very messy. For a pervasive such obtuseness, see C. Travis, 'Order out of Messes', Mind, vol.104, (1995)

<sup>x</sup> See his "Nice Derangement of Epitaphs" in Perspectives on Truth and Interpretation, ed., E. Lepore (Blackwell, 1985)

---

<sup>xi</sup> See P. Boghossian “The Rule Following Considerations.” Mind, vol.98, (1989)

<sup>xii</sup> I don’t mean to say that that intention does not have even an indirect relevance to meaning. How exactly it does is a longer story which I won’t tell here, but in the telling it would become clear that it would give no support to the sort of normativity that Kripke and others think attach to the meanings of words like ‘fly’, i.e., to lexical meaning.

<sup>xiii</sup> To say “I am going towndown’ to convey that I am going downtown certainly puts the hearer to strain so it does indeed violate this utilitarian norm. But it violates no norm that is intrinsic to meaning since there are no such intrinsic norms. Or so I claim.

<sup>xiv</sup> Perhaps the first to suggest the relevance of orthodox accounts of reference to questions of theory-change and meaning-change was Israel Scheffler in his Subjectivity and Science. Putnam makes the point too in Putnam, *ibid.* Fodor makes the point in various books and articles, but see especially his and E. Lepore’s Holism: A Shopper’s Guide, Blackwell (1992) See also Philip Kitcher, “Theories, Theorists, and Theoretical Change’ in Philosophical Review, vol.87, 1978.

<sup>xv</sup> Davidson “Radical Interpretation”, in Davidson, *ibid.* Evans, Varieties of Reference (Oxford University Press, 1980), and McDowell, “Singular Thought and the Extent of Inner Space” in J. McDowell and P.Pettit, Subject, Thought and Context, (Oxford, 1986)

<sup>xvi</sup> See his “Truth Rehabilitated” in Rorty and His Critics, ed. R. Brandom (Blackwell, forthcoming). Here Davidson says the concept of truth is objective in the sense that one could never know when any of our beliefs are true. I think there is no reason to tie the objectivity of truth to something like that idea. But the point for now is that if that is how he thinks of truth, then the notion of conditions of truth, which is supposed to give the meanings of sentences and the contents of our beliefs, is never going to meet the first Fregean constraint. See my “Is Truth a Goal of Inquiry” in the same volume as Davidson’s essay for more on this and other aspects of Davidson’s views on truth and meaning.

<sup>xvii</sup> As should have been clear from footnote 1, I am departing from custom a bit because I do not restrict the term “concept” to talk just of what is expressed by the use of predicates but extend it to talk also of what is expressed by the use of names. Thus I often talk of Pierre’s ‘London’-concept just as I might talk of someone’s ‘water’-concept.

---

<sup>xviii</sup> Its worth treading carefully because there are familiar issues involved here about substitution of extensionally equivalent terms in belief contexts. Ebbs presents (1)-(3) as being about everyday practice of content-ascription. It is not obvious whether everyday practice does or does not obey the familiar caution that we cannot substitute extensionally equivalent terms in these opaque contexts unless the speaker was aware that the terms or concepts are extensionally equivalent. If everyday practice were to attribute contents in a way that was sensitive to an agent's conceptions of things (conceptions of substances, diseases, stars, cities), then it would of course obey the caution since, for example, it would insist that we not assume that a medically ignorant man always had the same 'arthritis'-concept that the medically knowledgeable man has. But then to the extent that it does obey the caution in this way, it was my point that even everyday practice of content-ascription would *abandon* orthodox externalism, for there is nothing left of orthodox externalism, nothing left of the idea that extension characterizes the concepts which compose contents, if such substitution is not allowed in content-ascription. In the book I had raised this difficulty for the orthodox externalist and then pointed out that the orthodox externalist could take the desperate measure of saying that one cannot substitute anything whatsoever in these contexts. This is the view I had called the mysterious view of concepts, which Ebbs upholds and about which I say something below.

<sup>xix</sup> There is a very bad argument abroad for why there must be a distinction between concepts and conceptions, why there must be concepts over and above conceptions. It is that if some concepts have explications (i.e., if some concepts are to be characterized in terms of conceptions), others cannot be because otherwise there will be a circle. I address and answer this argument in the book on p. 46. But, in any case, it seems to me to be no measurable improvement to say that we should get rid of a circularity (assuming for the moment it is really there, which it is not) and introduce instead a total blank and mystery.

<sup>xx</sup> I don't mean to deny, of course, that this methodological canon about the revisability of the explananda requires us to then say a great deal about what gives structure to inquiry and to theory-construction, whether in the study of meaning and intentionality or of anything else; but whatever else we do say about that is not going to cancel this methodological canon.

<sup>xxi</sup> See pp.188-190.

<sup>xxii</sup> See my paper "Self-Knowledge and Resentment" in On Knowing One's Own Mind, eds., C. McDonald, C.Wright, and B.Smith, (Oxford University Press, 1998)

---

<sup>xxiii</sup> For references to Fodor and Searle on these issues, see Fodor and Lepore, *ibid.*, and Searle, among other things, [The Rediscovery of Mind](#) (MIT Press, 1992). By the way, I don't at all mean to suggest that Searle and Fodor share much apart from their rejection of a view which stresses rationality in action for not getting right the intrinsic nature of intentionality and therefore failing to be genuinely realist about intentional states. Their respective positions are very different apart from agreement on this. Moreover even their reasons for rejecting the above view come from different sources.

<sup>xxiv</sup> Oxford, 1986. See my discussion of this work and the general issue in [Belief and Meaning](#), ch.5.

<sup>xxv</sup> See Barry Stroud, "Transcendental Arguments", [Journal of Philosophy](#), vol.65, 1968. Actually Stroud's word to describe the replaced premise is not 'anti-metaphysical realist' or any earlier equivalent of that. It is 'verificationist'; and that description is thoroughly misleading. There is nothing verificationist about the assumptions involved in most of the standard transcendental arguments against the sceptic, they are merely not metaphysically realist assumptions. Verificationism is a very specific doctrine and so the term should not be cavalierly applied to all denials of metaphysical realism. Nagel, in the work mentioned above, also succumbs to the tendency of describing these denials with polemically charged and misleading terms, such as 'idealism'.