

Chapter three: Individualism, Non-Individualism, and Anti-Individualism

While it is generally accepted that the ascriptional practices Kripke appeals to (especially those involving proper names discussed in chapter one) cannot be excluded from any acceptable equilibrium, the status of those ascriptional practices Burge appeals to is considerably less well entrenched.¹⁶² The purpose of this chapter will be to argue that our practice of interpreting people in terms of the ‘accepted’ usage in their community can, indeed should, be incorporated into equilibrium for our semantic terms.

One set of terminological points that should be made clear from the start is the following. The term “individualism” will be here associated with the belief that an individual whose usage is known to stray from the social norm should *never* be interpreted according to the usage of his community; “anti-individualism” will be associated with the position *contrary* to individualism, i.e., the belief that such a speaker should *always* be interpreted in terms of the accepted usage in his community, and “non-individualism” will be associated with the position *contradictory* to individualism, i.e., the belief that an individual whose usage strays from the social norm should at least *sometimes* be interpreted in terms of the accepted usage in his community.¹⁶³ It will be argued below that incorporating the practices Burge appeals to requires only that one adopt a *non-individualistic* position; it does not require adopting the stronger *anti-individualistic* position.¹⁶⁴

¹⁶² The practices Putnam appeals to occupy something of an intermediate position, and will not be discussed directly here. Most who are willing to accept Burge’s intuitions will also be willing to accept Putnam’s, and any argument that supported the practices Burge and Wilson draw to our attention (which the following two chapters will attempt to do) will also constitute an argument in favor of the Putnam practices.

¹⁶³ This terminology is, once again, taken from Katarzyna Paprzycka’s discussion of a similar taxonomy of behavioral explanations. (Paprzycka forthcoming.)

¹⁶⁴ Burge himself is best understood as defending a non-individualistic rather than an anti-individualistic position, and critics of Burge such as Davidson and Bilgrami are also, ultimately, best understood as defending non-individualistic rather than individualistic positions (though they are clearly committed to *methodological* individualism in way a in which Burge is not). Still, while Davidson’s and Bilgrami’s considered positions are not actually individualistic, they can still be considered “defenders of individualism” since they do, in fact, go to considerable length to defend it (at least in part because they take Burge to be defending an anti-individualistic position). For a discussion of these issues, see appendix D.

There are, of course, other criteria by which a position could be characterized as “individualistic,” and the position defended in this chapter, while non-individualistic in the terms described above (hereafter “ascriptional” individualism”), is “individualistic” in certain other important respects. In particular, while this chapter argues against ascriptional individualism, it does endorse “methodological” individualism. Methodological individualism requires only that, when interpreting an individual speaker, we count ‘external’ factors (such as, say, the usage of his surrounding community) as relevant in determining what he means only if the interpreter himself takes them to be so. Such a position is, in many respects, at least as interesting as ascriptional individualism, but when we are concerned with whether the actual ascriptional practices that Burge focuses upon can be accommodated into an equilibrium, ascriptional individualism is what is at issue. This is, of course, because it is only ascriptional individualism that is necessarily incompatible with the acceptability of such ascriptional practices. (Methodological and ascriptional individualism are often run together, and they are discussed and distinguished in greater detail in the appendix D.)

This chapter will first investigate which of individualism, non-individualism and anti-individualism is best supported by our ascriptional practices, and then consider various arguments for taking these ascriptional practices more or less seriously. It will then turn to our general beliefs to see if and how they can be accommodated by the three positions. It will turn out that the non-individualistic alternative is not only best supported by our ascriptional practices, but it is also compatible with most of the general beliefs frequently taken to be incompatible with it.

1 The Ascriptional Practices

1.1 The relation between the Burge and the Kripke practices

As stated before, Kripke’s intuitions about proper name reference in the cases described in chapter one are almost universally shared. Even defenders of description

theories tend to accept the intuitions -- they just try to argue that they can be explained within the descriptive framework. On the other hand, the intuitions that Burge appeals to are considerably less secure, and the first reaction of many readers (including myself) when presented with Burge's thought experiments was simply to deny that, say, Bert meant what most other people did by "arthritis."

Kripke's work on proper names can thus be seen as having the potential to shore up a more generalized externalism in two ways. First of all, it can 'prime' our intuitions in a way that makes the conclusion Burge draws from his thought experiments seem more appealing.¹⁶⁵ Secondly, and much more significantly, if it turned out that the best account of the phenomena Kripke points to entailed that Burge's conclusions also followed, then Burge's conclusions would share the intuitively secure footing of Kripke's work. This chapter will argue that this is just what happens: the best explanation of the Kripke ascriptions allows one to endorse the Burge ascriptions as well.

Whether or not an equilibrium for our semantic terms would incorporate the types of ascriptions Burge appeals to¹⁶⁶ may depend, then, to a considerable extent upon whether or not these ascriptions are best understood as different aspects of the very same practices that Kripke appeals to. The anti-/non-individualists will thus have a considerable interest in showing that the ascriptional practices appealed to in the Burge cases are really of a piece with those appealed to in the Kripke cases. The individualist, on the other hand, will be concerned with describing the ascriptions that Kripke and Burge appeal to as instances of two distinct practices rather than as different aspects of the same practice. If he succeeds in doing this, the individualist can say that while we may have to accept

¹⁶⁵ As I said in chapter one, I doubt that the intuitions behind Burge's thought experiments would carry the force they do if the general beliefs they cut against hadn't already been shaken by Putnam's and Kripke's work. *Naming and Necessity* could, in some sense, be understood as being at least part of a paradigm shift in philosophical thought about language. After this shift, appeals to the sort of thought experiment Burge relies upon had already been entrenched in the profession, and had they not already been so, the comparative infirmity of the intuitions Burge appeals to might have been enough to dismiss them.

¹⁶⁶ Which, we must remember, are not the same as our reactions to thought experiments. (See chapter 1, sect. 8.)

Kripke's deeply entrenched set of ascriptional practices, Burge appeals to a different set that can be left out of any equilibrium.

While engaging in such a project, the individualist could appeal to in the widespread assumption that, since the works of Kripke and Burge highlight different things about the contents of our words and thoughts,¹⁶⁷ there should be two (possibly complementary, possibly rival) 'externalist' theories of content: 'causal/historical' theories and 'social context' theories.¹⁶⁸

The anti-/non-individualist would obviously like to show that the line of thought above is mistaken, and that whatever differences there are between Kripke and Burge, they can be understood as relating largely to the types of examples focused on, not to the general practices they draw upon. The fact that they highlight different aspects of the *same* practice has been obscured, the anti-/non-individualist will argue, by a number of factors. The first of these is that much of the excitement first generated by Kripke's work related to its prospects for providing a physicalistically respectable account of reference and rehabilitating certain types of essentialism.¹⁶⁹ Both of these projects seem quite alien to anything that would come out of Burge's work, and the focus on them caused the social, historical, and genetic aspects of Kripke's account to be ignored. Secondly, Burge himself initially presented the phenomena in "Individualism and the Mental" as being

¹⁶⁷ Namely, Kripke's work (along with Putnam's) illustrates that content is in part determined by our *physical environment* while Burge makes a structurally similar point relating to our *social context*.

¹⁶⁸ For a recent argument for the distinctness of Burge's and Kripke's positions, see Donnellan 1990. This assumption, though not often explicitly stated, is often implicit in criticism and expositions of Kripke, Putnam and Burge. Criticisms of Kripke and Putnam based on the assumption that they rely on metaphysically natural kinds presuppose that they provide a different sort of theory from Burge, who clearly requires no such assumption. Further, expositions of Burge according to which he argues that a society's present application of a term determines its meaning, and thus cannot account for Putnam's intuitions that we meant something different by "water" than our Twin-Earth counterparts even *before* water's chemical structure was discovered (Pettit & McDowell 1986b, 7-10), also presuppose that two different types of external contribution are being pointed to. Of course, this is not to deny that everyone recognizes the two points are of the same general *type* and that all the writers involved are externalists.

¹⁶⁹ See, for instance, Field 1972, Devitt 1980, and Salmon 1982. Putnam's work (especially Putnam 1975a), inspired similar hopes, though it seems quite clear that Putnam does not share them (see, for instance, Putnam 1981c, 1981d, 1988).

unaccountable for by the picture presented by Kripke,¹⁷⁰ thus making his work appear to present a rival paradigm rather than developing an earlier one. One of the main reasons, then, for thinking that Kripke and Burge are talking about different practices comes from the tendency to view Kripke as offering a plan for a non-intentional reduction of reference according to which causal relations and the ‘metaphysical essences’ of various kinds determine what we mean. Whatever Burge was up to, it certainly wasn’t that.¹⁷¹

As a result, the anti-/non-individualist will be concerned with showing, then, that Kripke did not, in fact, present a theory in which causation and the world’s essential structure could alone determine the meanings of our words.¹⁷² The anti-/non-individualist will thus stress that Kripke is not only notoriously cagey about referring to his work as a “theory” at all,¹⁷³ but he also explicitly states that his account does not eliminate reference from the explanans.¹⁷⁴ Kripke is perfectly willing to call his ‘picture’ a

¹⁷⁰ Or, for that matter, Putnam. While Burge admits that there are points of contact between his picture and Putnam’s (Burge 1979, 73, 117, 118), he still believes that the argument extracted from Putnam’s paper “is narrower in scope,” and “seems to work only for natural kind terms and close relatives” (Burge 1979, 118). He claims that purely causal accounts cannot account for the examples he discusses for the following reason.

The problem is unaffected by the suggestion that we specify input and output in terms of causal relations to particular objects or stuffs in the subject’s physical environment. Such specifications may be thought to help with some examples based on indexicality or psychological success verbs, and perhaps in certain arguments concerning natural kind terms (though even in these cases I think that one will be forced to appeal to intentional language). . . . But this sort of suggestion has no easy application to our argument. For the relevant causal relations between the subject and the physical environment to which his terms apply -- where such relations are non-intentionally specified -- were among the elements held constant while the subject’s beliefs and thoughts varied. (Burge 1979, 107.)

These claims are, in some sense, unobjectionable, but Burge would be very much mistaken if he saw them as applying to either Kripke or Putnam (rather than to their purely causal/reductive reconstructions). Recently Burge has characterized Kripke’s work as showing that reference depends on “a mix of causal and intentional elements and includes a person’s reliance on others to fix the referent” (Burge 1992, 24), and he seems to have mellowed a bit on Putnam as well (though this may be because he sees Putnam as having changed his views (Burge 1992, 47)).

¹⁷¹ Another factor that made the two sets of practices seem different was the unfortunate tendency to understand Burge as defending the type of quasi-communitarian picture of linguistic norms that would seem to have more in common with the reading of Wittgenstein purportedly in Kripke 1982 than it does with Kripke’s own views. (I have doubts about this reading of Kripke 1982 as well, but one can see instances of such quasi-communitarian readings of both Burge and Kripke’s Wittgenstein in, for instance, Bilgrami 1992a.)

¹⁷² Reasons for thinking that neither Kripke nor Putnam should be understood this way are discussed in appendix A.

¹⁷³ Kripke 1972, 93.

¹⁷⁴ Kripke 1972, 94, 97.

“historical” one, and it is only later writers who have put the emphasis on causality. Kripke’s picture of reference involves two parts: the ‘baptism’ at which the word is first associated with an individual or kind, and the transmission through which the meaning is preserved through subsequent use.¹⁷⁵ While both aspects of this picture can be, and have been, given causal/non-intentional interpretations,¹⁷⁶ such readings misrepresent Kripke’s text and ignore the possibility that Kripke’s work could provide the foundation for a more general externalist theory.

Not only can Kripke’s account of baptism be generalized beyond proper names and natural kind terms,¹⁷⁷ but his account of transmission can also be so extended. Kripke argues that a speaker refers to whatever he names “by virtue of his membership in a community which passed the name on from link to link,”¹⁷⁸ and there is nothing in this story that suggests that this linking is only possible with proper names and natural kind terms. All that is required of such a linking is that “the speaker intends to use the name in the same way as it was transmitted to him.”¹⁷⁹ Indeed, there is little reason to think we would have such intentions for proper names and natural kind terms but not for terms

¹⁷⁵ “A rough statement of a theory might be the following: an initial ‘baptism’ takes place. Here the object may be named by a description. When the name ‘is passed from link to link,’ the receiver of the name must, I think, intend when he learns it to use it with the same reference as the man from whom he heard it.” (Kripke 1972, 96.)

¹⁷⁶ See Devitt 1980, x. In the case of the baptism, a purely causal reconstrual would involve either the person or the initial sample (and its ‘essential’ structure) entering into causal relations with the baptizer (Devitt 1980, 27-31, 138.). However, such a reading would have trouble making sense of Kripke’s claims that “there need not always be an identifiable initial baptism” (Kripke 1972, 162). Furthermore, words whose meaning is the product of a reference shift could all be understood as, in some sense, lacking an initial baptism.

¹⁷⁷ Once again, for a discussion of this, see appendix A.

¹⁷⁸ Kripke 1972, 91. This appeal to the speaker’s being embedded in a social/historical environment is by no means isolated (see, for instance, Kripke 1972, 94, 95).

¹⁷⁹ Kripke 1972, 163. Devitt considers the reliance on these reference-preserving intentions to be a weakness of Kripke’s account, and tries to give an purely causal replacement of them (Devitt 1980, 138). Of course this appeal to reference-preserving intention may seem too intentional and voluntaristic, and it may be misleading to say that when we use a name or kind term, we use it with the intention that it refer to whoever or whatever those from we learnt it referred to by it. However, much the same result can come from the speaker’s having the belief that he is using the name or kind term in the same way as it was transmitted to him, and this belief, even if not explicitly formulated, can be seen as a consequence of the speaker’s belief, or even presupposition, that the term is part of a shared common language that he is speaking. For a discussion of how one can get at what Kripke wants here in a less voluntaristic fashion, see appendix F.

such as “arthritis” or “contract.”¹⁸⁰ The commitment the speaker takes on through this intention is not the type of commitment that he could only form with proper names and kind terms, and so if such intentions and commitments can determine what our names and kind terms refer to, they should be able to do so for our other terms as well.

This is the source of the following general belief about meaning:

If the speaker intends to use a term in the same way as it was transmitted to him then the object referred to by the speaker with the term is the same as the object referred to by the transmitter with the term.

As argued above, the acceptance of this belief (along with some plausible assumptions about the speaker’s intentions) leads almost immediately to the acceptance of the ascriptional practices Burge appeals to. As a result, the individualist will have some reason to want to reject such a belief, though it is hard to see how to reject it without also giving up the Kripke practices. The story that is in play about the transmission of names and kind terms already allows that a speaker can use them at least partially in virtue of his assuming that the term in question belongs to a language shared between him and the other users of the term that he has encountered.¹⁸¹ Once this reference-preserving mechanism is in play, there seems to be no reason for it not to be extended to terms in the language other than proper names or natural kind terms.

For instance, if John means *arthritis* by “arthritis,” and Bert picks up the term from John intending to use it with the same reference as him, Bert should mean arthritis by “arthritis” as well. The fact that Bert believes that one can get arthritis in one’s thigh is no more of a problem than the fact that our belief that Thales said “all is water” may turn out to be false. Since the process involved does not require that our beliefs about the referent be true in order for us to talk about it (that Thales said “all is water” might be all a speaker (explicitly) believes about him), Bert’s words and thoughts are connected to

¹⁸⁰ Especially since a given speaker or community may not be aware of whether a given term of his picks out a ‘natural’ kind.

¹⁸¹ This assumption need not, of course, be explicitly made. It may be the sort of implicit ‘assumption’ that forms part of what Searle describes as the “Background” against which our intentional states are found (Searle 1983). For a fuller discussion of these issues, see appendix F.

arthritis in spite of his mistaken beliefs about where it can occur. It should go without saying that this Kripkean story will predict that, in a counterfactual situation where the community that Bert has been brought up in means something different by “arthritis,” Bert will as well. It should also be noted that this application of Kripke’s picture explains Burge’s intuitions by making an appeal only to the account of transmission. The more problematic and controversial picture of the baptism need not be brought into play, since it is simply taken as *given* in Burge’s thought experiments that the members of the surrounding communities mean what they do by their terms.

It seems, then, that the practices appealed to by Kripke are best understood as encompassing the Burge ascriptions as well. Attempts to describe them as distinct (one practice tying content purely to causal history, another purely to social context) fail to do justice to the practices themselves. They also fail to do justice to the philosophers in question, as Kripke’s picture comes out as unrealistically essentialist, and Burge’s as overly anti-individualistic and unrealistically communitarian.

1.2 The status of the Burge practices

In addition to challenging the connection between the Kripke and the Burge practices, the individualist can also challenge the entrenchment of the Burge practices directly. There are at least three ways to do this. The first is to argue that we don’t need to incorporate such attributions into an equilibrium because we do not, in fact, make attributions this way. The second is to argue that, even though we do make such attributions, they can be ‘explained away’ pragmatically and do not reflect our considered judgments. The third is to argue that even if such ascriptions do reflect our pre-theoretical judgments, such judgments are of little theoretical significance. It will be argued below that all three approaches are ultimately unconvincing.

1.21 On the denial of the practice

If we did not make the kind of attributions Burge claims we do, there would clearly be no pressure to incorporate such attributions into any equilibrium, since they would turn out not to be entrenched at all. One writer who has frequently argued that our ascriptional practices themselves actually favor individualistic content attributions is Davidson. He has several reasons for thinking this. The first of these is that there are cases where we clearly do seem to make content attributions individualistically, and it would certainly strengthen the individualist's position if he could show that our ascriptional practices do not uniformly support anti-/non-individualism the way they are often represented as doing. Davidson focuses on what he takes to be such cases quite frequently.

If you see a ketch sailing by and your companion says, 'Look at the handsome yawl,' you may be faced with a problem of interpretation. One natural possibility is that your friend has mistaken a ketch for a yawl, and has formed a false belief. But if his vision is good and his line of sight favorable it is even more plausible that he does not use the word 'yawl' quite as you do, and he has made no mistake at all about the position of the jigger on the passing yacht. We do this sort of off the cuff interpretation all the time, deciding in favor of reinterpretation of words in order to preserve a reasonable theory of belief. As philosophers we are peculiarly tolerant of systematic malapropism, and practiced at interpreting the result. The process is that of constructing a viable theory of belief and meaning from sentences held true.¹⁸²

Davidson's choice here is not a particularly happy one. With respect to the ketch/yawl example, one can freely admit that the speaker "does not use the word 'yawl' quite as you do," while still insisting that the speaker has made a (non-perceptual) mistake and does mean by "yawl" what the interpreter does. Indeed, the question of whether difference in 'use' in this restricted sense¹⁸³ must involve a corresponding difference of meaning is precisely what is at issue between the individualist and the non-individualist. Matters are not helped by the fact that the case itself is underdescribed. If the speaker is, like myself, generally ignorant of sailing terminology and sailboat construction, I would be inclined think that he simply cannot tell ketches and yawls apart and has falsely called a ketch a

¹⁸² Davidson 1974b, 196.

¹⁸³ That is, in the sense corresponding to how the word is actually applied by the speaker, as opposed to the more 'generous' notion of "use" in which we can be said to use the word "arthritis" to talk about arthritis even if we occasionally misapply the term.

yawl. On the other hand, if the speaker is quite familiar with the two types of boat and *systematically* calls ketches “yawls” and yawls “ketches” (and perhaps is able to explain that the two types of boats are distinguished by the position of the jigger etc.), then the claim that he cannot tell the two types of boat apart seems clearly false and the claim that he means *ketch* by “yawl” seems considerably more plausible. If non-individualism is true and what sorts of ascriptions we make must be decided on a case-by-case basis, it may frequently turn out that a disputed example is simply underdescribed.

1.211 Slips, malapropisms and misunderstandings

Nevertheless, even though I do not share Davidson’s intuitions in this *particular* case, Davidson’s invocation of malapropisms should make clear the types of case he has in mind. In particular, there seem to be three types of case where we tend not to understand the interpretee, and the interpretee does not understand himself, in terms of the meanings conventionally associated with his words:¹⁸⁴

(1) Slips of the tongue, slips of the pen, spoonerisms and typos.

All of these ‘performance errors’ are pretty clear cases where we do not interpret someone according to the meanings conventionally associated with his words. If, after I have looked at Bill’s diary, I type “I was just looking at Bill’s dairy,”¹⁸⁵ this is not taken to be an expression of a false belief that I was looking at Bill’s *dairy*, but rather as a clumsy expression of a true belief that I was looking at Bill’s *diary*. “Dairy” is here either interpreted as meaning *diary*, or it is interpreted as meaning dairy, but no judgement is attributed corresponding to the assertion made.¹⁸⁶

¹⁸⁴ Burge discusses these three types of case as instances where we do tend to reinterpret the speaker. (Burge 1979, 90.)

¹⁸⁵ The sort of mistake that, as any reader of earlier drafts of this knows, I make all too frequently.

¹⁸⁶ In terms of the discussion in chapter two, the expressive constraint does not hold for such performance errors.

(2) Malapropisms, systematic mispronunciations and spelling mistakes

Malapropisms such as “we need a few laughs to break up the monogamy” are also cases where we are not inclined to interpret a speaker in terms of the meanings conventionally associated with his words. For instance, until quite recently, I often confused “eczema” and “emphysema.” As a result, I would say things such as “Peter’s smoking has made his eczema much worse.” In such cases, I was not understood, and did not understand myself, as making the false claim that Peter had a skin problem that was exacerbated by his smoking. I was simply understood as meaning *emphysema* by “eczema.”¹⁸⁷ As with slips of the tongue and mispronunciations, we are inclined to say “what he meant to say was ...” rather than understanding the speaker as having a false belief.¹⁸⁸

(3) Radical misunderstandings.

If his misunderstanding is radical enough, there is also a tendency to reinterpret the speaker so that he is not actually speaking and thinking falsely about the objects that his words are conventionally associated with. For instance, if Bert came to believe that “arthritis” was a retirement community in southern Florida, we would probably be disinclined to treat his “arthritis” utterances as being about *arthritis*. The more radical the speaker’s misunderstanding is, the greater percentage of his beliefs will have to be given up in order to accommodate the beliefs tying his use to communal standards. As a result, the more radical the speaker’s misunderstanding, the more likely it will be that the equilibrium for his total belief set will not include these more socially directed beliefs.

¹⁸⁷ Idioms provide many cases like this. People new to a language frequently make minor but often amusing errors with their idioms. I know of a woman who once said at work “I came in late and the boss ate me out good”; she is understood by everyone as truly saying and believing that she was severely criticized verbally (“chewed out”), not as having a false belief about her sex life.

¹⁸⁸ We are more likely to say “by ‘epitaph’ I meant ‘epithet’” or “by ‘orangutan’ I meant ‘orange Tang’” than we are to say “by ‘arthritis’ I meant ‘something very much like arthritis that one can also get in one’s thigh.’”

1.212 Individualistic vs. individualized interpretations

It is, of course, no coincidence that the sorts of cases where a more individualistic interpretation seems tempting are precisely those that caused trouble for the ‘meaning-theoretic’ accounts of the expressive constraint discussed in chapter two. However, while such cases may force us to give up purely meaning-theoretic accounts of belief, they do not necessarily mean we should give up non-individualism or even anti-individualism. While such cases may show that the interpretation may need to be *individualized*, it need not follow that the interpretation idiosyncratic speaker must be *individualistic* in the sense of being *independent* of standard usage. (Indeed, given obvious differences of pronunciation etc. our interpretations of others will have to be individualized to some extent at at least that level.)

One could admit that the speaker’s idiolect remains primary when it comes to deciding which terms are assigned to which concepts, and still claim that a speaker’s concepts are determined by the usage of the community. We each speak our own idiolect, and while these idiolects may be different for each person (and each person at each time), they remain dependent upon the shared linguistic norms. Malapropisms and spoonerisms provide fairly clear examples of this. When Mrs. Malaprop says (while pointing to an arrangement of epithets) “What a nice derangement of epitaphs” one interprets her as meaning *epithet* by “epitaph.” This interpretation works only for Mrs. Malaprop and, in this sense, there is an idiolect that only she speaks. Nevertheless, the account does not treat her usage as unconstrained by social norms. On such a theory “epitaph” is true not of just those things Mrs. Malaprop would call “epitaphs,” but rather of those things that are, in fact, *epithets* -- where what counts as an epithet is tied to the usage of the community. One cannot just read the meanings conventionally associated with the terms in the language into the speaker’s utterances, but even when we come up with a theory uniquely suited to the speaker, such a theory may still make use of socially

determined concepts.¹⁸⁹ Cases of the first two types (slips and malapropisms) will usually be of this nature.

Still, while the first two types of case (slips and malapropisms) may ultimately not be best understood in completely individualistic terms, it seems quite possible cases of the third sort (radical misunderstanding) may favor a more individualistic understanding. A speaker may understand a term in an extremely idiosyncratic way that cannot be mapped on to any socially determined concepts. In such cases, it may be best to interpret the speaker as using a concept that is not in the social stock. Though even here, it isn't clear that the speaker is being understood individualistically. In particular, it may be that the speaker has a novel concept partially in virtue of combining other concepts in an idiosyncratic way, but the concepts so combined may themselves be socially determined. I may be the only person to mean *beer from Ireland* by "porter," but what I mean by "beer" and by "Ireland" (and thus what I ultimately mean by "porter") may still be determined by social use.

Nevertheless, while the fact that there may be *some* cases that must be understood this way causes serious problems for anti-individualism, it does not really support individualism, since individualism requires that *all* cases be of this sort. As a result, the fact that we may, at times, interpret others 'individualistically' suggests that our ascriptional practices favor non-individualism over anti-individualism, but not that they favor individualism over non-individualism.

¹⁸⁹ Indeed, the point could be seen as an extension of Evans' discussion of linguistic "counters" below:

When we wish to establish what a person is saying in using certain words, we must get clear exactly which linguistic counter, so to speak, the speaker is putting forward. Here the speaker's intentions are indeed paramount. But once it is clear which linguistic counter he is putting forward, the content of what he says is determined by the significance which that counter has in the game. (Evans 1982, 68-9.)

Not only can we identify the counter the speaker intends to play in spite of the fact that he may have various mistaken ideas about the counter's value (arthritis-cases), but we can also identify the counter he has in mind in spite of his using a word that is conventionally associated with a different counter (malapropisms).

1.213 Davidson on “arthritis”

Davidson, however, claims that our ascriptional practices not only favor non-individualism over anti-individualism, but also favor individualism over non-individualism. He is thus frequently¹⁹⁰ concerned with arguing that we do, in fact, always interpret each other individually. As a result of this, Davidson suggests that when Burge claims that we would say that Bert means *arthritis* by “arthritis,” he is describing his own theory-driven intuitions rather than our deeply entrenched ascriptional practices.

Davidson claims not to be impressed by Burge’s arguments because he has “a general distrust of thought experiments that pretend to reveal what we would say under conditions that in fact never arise.”¹⁹¹ However, it is hardly fair of Davidson to argue that Burge’s arguments rely on our intuitions about “what we would say under conditions that in fact never arise.” While such a charge could perhaps justly be made against arguments relying heavily on, say, Twin-Earth cases, examples such as Bert’s use of “arthritis” are of a sort that we run into frequently in everyday practice. I was recently with a group engaged in a conversation when it became clear that one speaker thought that “vests” could have sleeves. The speaker had a number of true beliefs about vests (that they were items of clothing, that I was wearing one at the time, that they form an essential part of three-piece suits etc.) and could recognize most vests when he saw them, but would also be willing to apply the term to ‘sleeved vests’ and other “jackets that weren’t quite up to snuff.” Nobody (including the speaker himself) took the speaker to mean anything different than usual by “vest.” He was simply understood as having a false belief about vests. Similar cases come up all the time in everyday discourse (John thinks that any coffee with steamed milk is a cappuccino, Susan thinks that sherbet is a

¹⁹⁰ But by no means always (see appendix D).

¹⁹¹ Davidson 1991b, 197. It should be noted in this context that Davidson’s discussion of the “Swampman” is not a thought experiment of this sort. He is not trying to bolster a theory by appealing to our intuitions in this case, but using a case to illustrate the consequences of a theory he has already defended. Indeed, the example is effective precisely because what he takes to be the correct thing to say about the Swampman is *counterintuitive*.

type of ice cream, etc.) and there is no reason to think that the purported phenomena rest on intuitions about various types of ‘science-fictionish’ thought experiments of the sort familiar from, say, debates about personal identity.

However, Davidson does not merely argue that the types of cases Burge appeals to “never arise.” He also doubts that someone “unspoiled by philosophy” would make ascriptions the way Burge says they would.¹⁹² Davidson argues that, in precisely the type of case Burge appeals to, it is simply “false that our intuitions speak strongly in favor of understanding and interpreting an agent’s speech and thoughts in terms of what others would mean by the same words.”¹⁹³ Rather, he claims that our actual practice favors the individualistic interpretation. He gives an example of and motivation for such ‘correct’ interpretive practices in this passage:

Suppose that I, who think the word ‘arthritis’ applies to inflammation of the joints only if caused by calcium deposits, and my friend Arthur, who knows better, both sincerely utter to Smith the words “Carl has arthritis”... If Smith (unspoiled by philosophy) reports to still another party ... that Arthur and I both have said, and believe, that Carl has arthritis, he may actively mislead *his* hearer. If this danger were to arise, Smith, alert to the facts ... would add something like “but Davidson thinks arthritis must be caused by calcium deposits.” The need to make this addition I take to show that the simple attribution was not quite right; there was a relevant difference in the thoughts Arthur and I expressed when we said “Carl has arthritis.”¹⁹⁴

According to Davidson, then, our ascriptional practices, even in the sorts of cases Burge focuses on, will actually be individualistic.

However, with respect to the claim that such unqualified reports could mislead a speaker, it is, of course, open to the anti-/non-individualist to respond that their misleadingness is not incompatible with their truth. Such reports truly convey the thought expressed, but can give a misleading impression of the subjects *total* state of mind. If I tell someone who needs to go downtown that I have a car, I will have misled him if I do not also tell him that it is broken. The fact that such unqualified reports are misleading can be attributed to the pragmatics of the conversational situation rather than the content of what is said. Given that I know that the person wants to go downtown, the

¹⁹² Davidson 1987, 449.

¹⁹³ Davidson 1991b, 197.

¹⁹⁴ Davidson 1987, 449.

speaker will assume that I would not have mentioned the car if it was non-functional. If we are told that someone has a car, we will assume that it is in working order unless told otherwise. In much the same way, since most speakers know that arthritis is a disease of the joints, a hearer will assume that the person with an arthritis-belief will have such knowledge unless he is informed otherwise. As a result, one can mislead such a speaker by making the attribution without the qualification.

The extent to which such qualifications are necessary may be part of what determines the desirability of ascribing a socially determined concept. For instance, if Davidson believes that arthritis is caused by calcium deposits, but that is his only false belief about arthritis, then most of his behavior relating to arthritis (his purchase of arthritis medicine, his belief that he has it in his knee, his expression of sympathy to a friend who claims to have it, etc.) can all be explained with straight attributions, and additional qualifications about his beliefs about its cause will not be necessary. Of course, some of his arthritis behavior might need the added qualification (his cutting down his milk consumption because he believes he is getting arthritis), but this will probably be only a small fraction of it. On the other hand, for someone who believes that an orangutan is a fruit drink, it will turn out that just about *all* of his 'orangutan' behavior (his claiming that there is a cold pitcher of it in the fridge, his desire for some when he is thirsty, his claim that they can be purchased for \$1.25 at the corner store, etc.) will require that we mention the idiosyncratic belief, so the attribution of the 'standard' concept is considerably less useful.

Davidson is aware, of course, of such appeals to the pragmatics of the conversational situation, and he claims that they overlook "the extent to which the contents of one belief necessarily depend on the contents of others" and forgets that "thoughts are not independent atoms."¹⁹⁵ However, there seems to be no reason to think that the non-individualist (rather than the anti-individualist) needs to deny this. Such an appeal to

¹⁹⁵ Davidson 1987, 449.

pragmatics does not, as Davidson suggests, commit one to there being a “simple, rigid, rule for the correct attribution of a single thought.”¹⁹⁶ Unlike both the individualist and the anti-individualist, the non-individualist can take things on a case-by-case basis, so the claim that there should be no such rigid rule could be understood as favoring non-individualism.

Davidson here seems to be assuming that holism entails that any difference of belief will lead to a difference of meaning, and if he were right about this, holism probably would entail individualism.¹⁹⁷ There is, however, little reason to believe this.¹⁹⁸ If what a term refers to is determined by the equilibrium for all of the beliefs and applications associated with it, a change in a single belief will frequently make no difference as to what the equilibrium will be. One can thus admit that what a speaker means by his terms is a function of all of his beliefs, without being committed to the claim that any change in his beliefs would constitute a change in what he means. This would only follow if the function in question is the identity function, and while some versions of ‘conceptual role semantics’ may work with such a function (and thus bring with them the counterintuitive consequences frequently associated with holism), the equilibrium driven story from chapter one clearly has a more ‘normative’ function in mind. As a result, while the account of meaning in chapter one is holistic, it avoids the counterintuitive consequences frequently associated with holism by both critics like Fodor and supporters such as Davidson.

Non-individualistic accounts can take best advantage of the flexibility that accompanies the radical interpretation process by dealing with utterances on a case-by-case basis and thus taking to heart Davidsonian claims such as “there can be no simple, rigid, rule for the correct attribution of a single thought,”¹⁹⁹ and “when others think

¹⁹⁶ Davidson 1987, 449.

¹⁹⁷ This is, of course, an assumption he shares with many of holism’s harshest critics. (See, for instance, Fodor and Lepore 1992)

¹⁹⁸ See the discussion of holism in chapter one for a fuller discussion of what comes immediately below.

¹⁹⁹ Davidson 1987, 449.

differently from us, no general principle ... can force us to decide that our difference lies in our beliefs rather than our concepts.”²⁰⁰ While a Davidsonian position certainly could not be anti-individualistic, it need not be individualistic either. A non-individualistic position seems in keeping with both the spirit and much of the practice found in Davidson’s writings.²⁰¹

1.22 ‘Explaining away’ the practice

Thus far we have been looking at attempts to show that we do not, in fact, ascribe attitudes in the way that Burge suggests. There is, however, another line of attack that, while admitting that we do make ascriptions the way Burge claims that we do, tries to downplay the significance of this part of our practice.²⁰² Doing so will usually involve trying to ‘explain away’ phenomena that seem to give at least *prima facie* support to anti-/non-individualistic positions. Attempts to explain away such phenomena can be seen in discussions of both our belief ascriptions and our deferential practices. (The two are not unconnected, since our deferential practices are a reflection of the crucial case of our self-ascriptions.)

1.221 Belief ascription

Davidson, for instance, has argued that, even if Burge is right about what we would say about Bert (which, as noted above, Davidson doubts), such ascriptions can be understood as reflecting laziness on the interpreter’s part or pragmatic constraints on ‘off the cuff’ interpretation rather than actual truths about what he means by his terms. Given the number of comparative strangers we meet each day, there is no real possibility of

²⁰⁰ Davidson, 1974b, 197.

²⁰¹ Indeed, while Davidson tends to view himself as an individualist, his reasons for doing so tend to focus almost exclusively on problems with *anti*-individualistic positions, and Davidson’s own position (as distinct from his explicit pronouncements on the subject) is perhaps best understood as non-individualistic. (For a fuller discussion of this, see appendix D.)

²⁰² “We would, in everyday reporting of belief, say “Bert believes that he has arthritis in his thigh.” . . . It is quite open, however, whether this fact about everyday reports has anything to do with the issue before us. The issue before us is a philosophical one.” (Bilgrami 1992a, 74.)

engaging in full fledged radical interpretation with each of them. Because of this, we save time and effort by attributing ‘standard’ contents to their words.²⁰³ As Davidson puts it: “We do not have the time, patience, or opportunity to evolve a new theory of interpretation for each speaker ... But if his first words are, as we say, English, we are justified in assuming he has been exposed to linguistic conditioning similar to ours.”²⁰⁴ “When some of these assumptions prove false, we may alter the words we use to make the report, often in substantial ways. When nothing much hinges on it, we tend to choose the lazy way: we take someone at his word, even if this does not quite reflect some aspect of the speaker's thought or meaning.”²⁰⁵

However, this explanation of our ascriptional practices as being a natural response to the lack of evidence we usually have in conversation has problems accounting for the fact that we often stick with these ‘lazy’ ascriptions even after learning that the interpretee uses a term differently than we do. Indeed, if we do choose to flag any differences, we tend to flag them in terms of the speaker’s beliefs, not his meanings. Our tendency to interpret people non-individualistically cannot, then, be explained away by the claim that it is the product of time-constraints on off-the-cuff interpretations. Even if they had all the time in the world to interpret Bert, many interpreters (including, obviously, Tyler Burge) would still attribute to him the false belief that he has arthritis in his thigh.

Perhaps even more damaging to the claim that the non-individualistic elements in our ascriptional practices are produced by such pragmatic constraints is the fact that our *self-*

²⁰³ “This is not to deny the practical, as contrasted with the theoretical, importance of social conditioning. What common conditioning ensures is that we may, up to a point, assume that the same method of interpretation that we use for others, or that others use for us, will work for a new speaker. We do not have the time, patience, or opportunity to evolve a new theory of interpretation for each speaker, and what saves us is that from the moment someone unknown to us opens his mouth, we know an enormous amount about the sort of theory that will work for him -- or we know we know no such theory. But if his first words are, as we say, English, we are justified in assuming he has been exposed to linguistic conditioning similar to ours (we may even guess or know the differences). To buy a pipe, order a meal, or direct a taxi driver, we can go on this assumption. Until proven wrong; at which point we can revise our theory of what he means on the spur of the moment. . . . Knowledge of the conventions of language is thus a practical crutch to interpretation, a crutch which, under optimum conditions for communication, we can in the end throw away, and could in theory have done without from the start.” (Davidson 1982, 278-9.)

²⁰⁴ Davidson 1982, 278-9.

²⁰⁵ Davidson 1987, 449.

ascriptions are often made along non-individualist lines. It is very hard to reconcile the claim that non-individualistic content ascriptions are based upon our not taking the time to find out just what the interpretee really believes with the apparent fact that, in spite of our knowing what we believe, we make such non-individualistic ascriptions to ourselves. We may know perfectly well what we believe, and we still frequently see ourselves as having *misapplied* words and as having had *false* beliefs when we find that our usage deviates from the social norm. If we really self-ascribe contents non-individualistically, then the attempt to explain away non-individualistic content ascriptions by constraints on off-the-cuff interpretations seems broken-backed. The appeals to our ‘everyday ascriptions’ and deference behavior should thus be understood as part of a mutually supporting argument, with deference being crucially important not for its own sake, but rather as an indication of our self-ascriptions.²⁰⁶

It should not be surprising, then, that defenders of individualism have often tried to argue that we do not, in fact, make such self-ascriptions. To do so, they must give an account of what happens when we defer to standard usage that does not involve our understanding ourselves as having had false beliefs or as having misapplied our own terms. Such accounts are the topic of the next section.

1.222 Deference

Defenders of individualism will argue that deference is explained not by the fact that speakers understand themselves as having misapplied their terms, but rather by the fact that speakers know that, if they wish to be understood, they should use their words as others do. This allows deference to be given a pragmatic explanation.²⁰⁷ These more pragmatic and historical explanations of deference also open up an account of how we can be *mistaken* when our usage strays from the norm. We have a general intention to

²⁰⁶ Deference is thus connected to a crucial sub-case of our content attributions, and should not be understood as part of a separate, and perhaps unrelated, argument against individualism, as Bilgrami suggests Bilgrami 1992a, 78.

²⁰⁷ See, for instance, Bilgrami 1992a, 79-89, Davidson 1994a, 9.

speak as others do, and so we can be considered to have made a mistake each time we fail to realize this intention.²⁰⁸ These mistakes, however, consist in our failure to communicate effectively, not in our failure to understand how our words should be used.²⁰⁹

The gloss that individualists give to linguistic norms makes them much weaker than those proposed by Burge. In particular, they constitute a set of *hypothetical* rather than *categorical* imperatives for the speaker. That is to say, all of the ‘semantic norms’ in these cases are contingent upon the speaker having particular *desires* (usually having to do with communicating effectively). These desires do not, in themselves contribute to the term’s meaning. Bilgrami characterizes deference as a response to a norm “which necessarily has a hypothetical formulation,”²¹⁰ and Davidson claims that “any obligation we owe conformity is contingent on the desire to be understood.”²¹¹

However, the picture of the norms involved as being hypothetical in this fashion seems to misrepresent the phenomena we are dealing with. While there are undoubtedly cases where we alter our usage in order to promote ease of communication, such cases should not serve as a general model of what is happening when we ‘correct’ our usage.²¹² As a general psychological claim, saying that deference is a response to such a hypothetical norm relating to our desire to communicate is not especially plausible. When our usage is corrected, and we accept the correction, we usually understand ourselves as having made a mistake about the items in question (“I didn’t know that you could have oral contracts” etc.). We do not usually understand ourselves as having meant something different by the word, and as changing our usage just because the new usage is more communicatively efficient. People could come to see their mistakes this way

²⁰⁸ Bilgrami 1992, 110.

²⁰⁹ See Davidson 1994a, 9, Davidson 1986a, 434, Bilgrami 1992a, 110.

²¹⁰ Bilgrami 1992a, 110. He goes on immediately to specify the norm as saying “speak like others do, *if it pays to do so*,” or “more specifically,” “I ought to use words as others do, *if I want to be easily understood*.” See also Bilgrami 1992a, 111 for the claim that the norms involved are not only hypothetical, but also extrinsic (i.e.: make no references to, say, ‘English’).

²¹¹ Davidson 1994, 9.

²¹² Such a denial of generalizability is also found in Burge 1989, 184.

(Davidson and Bilgrami apparently do), but most speakers do not. When students get wrong answers on vocabulary tests, they are understood, both by their teachers and by their peers, as not knowing what the word in question means, not as speaking a different language that is poorly suited for doing well on tests. This understanding of our usage is not just a product of, and certainly extends beyond, the classroom.²¹³

The extent to which Davidson overinstrumentalizes our thoughts about language can be seen more clearly when we consider that Davidson portrays even small children as following pragmatic strategies when they learn their language: “The child babbles, and when it produces a sound like ‘table’ in the evident presence of a table, it is rewarded; the process is repeated and presently the child says ‘table’ in the presence of tables.”²¹⁴ Needless to say, there is no reason to think that the child needs any kind of reward system to encourage it to speak as others do in its environment. A child does not imitate the way those around her speak because she *wants* to be understood, or because she wants to speak, say, English properly. Imitative behavior begins at the phonetic level almost as soon as babbling begins, and speaking as others do seems to be inherently desirable for the language learner and is not dependent upon any other desires.²¹⁵ Indeed, the categorical nature of the imperative is such that it is almost misleading to say that the child *wants* to speak as others do at all; no such general desire needs to be present, each instance of imitative behavior is desirable in itself, not by reference to any more general desire. Even though it is the behavior of adults that we are interested in, this fact about infants illustrates, in a particularly pure form, a tendency to defer to surrounding usage that carries on throughout our lives.²¹⁶ This is not, of course, to deny that our actual need

²¹³ Which is not to say that our education plays no role in how this understanding manifests itself.

²¹⁴ Davidson 1992, 263.

²¹⁵ Of course things are a little more complicated than this, since issues of identification play some role in this process, but the type of sociolinguistic issues having to do with language learning and social and gender identification will not be explored here.

²¹⁶ Furthermore, even if one were to admit that the imperative to bring one’s use into line only presents itself as a hypothetical one, the norm of speaking, say, English correctly seems just as likely to be psychologically active as a norm related to ease in communication that makes no reference to a shared language. This becomes especially clear when we consider those who consciously stray from a type of usage that may now facilitate communication because of their belief that various new idioms do not

to communicate plays an important role in our understanding of deference behavior. However, while pragmatic factors may be involved in explaining deference behavior, allowing such factors to play some role is not incompatible with being a non-individualist. Pragmatic factors may shape which attitudes we have without these factors themselves being explicitly represented by us. Language would not be much use if we did not speak in roughly the same way, but these more ‘evolutionary’ considerations need not manifest themselves in any actual *desire* to speak as others so that we can be understood. Other sets of beliefs and desires can produce the same results just as, if not more, effectively.²¹⁷

For instance, one could explain deference by an appeal to a sub-personal mechanism that causes us discomfort whenever our behavior does not conform with that of those around us. We need not have any conscious *desire* to conform with our peers, we just find ourselves doing so. Our conscious beliefs and desires may just rationalize, rather than cause, the conformist behavior in question (in this case, deference to social usage). One could imagine the evolutionary use of such a mechanism, but it should be stressed that this does not count as a ‘pragmatic’ explanation of deference behavior of the sort that the individualist needs. If the pragmatic factors are not at the conscious level, and the agents actual beliefs and desires are still in terms of the agents having made a factual mistake about the objects in question rather than communal usage, then both the conditions of application and general beliefs associated with our words will favor the non-individualistic equilibrium.

It would, of course, be possible for a community to get together and simply decide to try to speak the same way, with certain members given the right to determine the

constitute part of “proper English.” Indeed, the widespread belief that languages are being ‘corrupted’ testifies to the presence of a conception of language which plays a normative role independently of its role in communication. This belief in the possibility of linguistic corruption may just be a reflection of our pre-theoretical conception of language, but, right or wrong, it is this pretheoretical conception that determines the motivational nature of the linguistic imperatives acting upon us.

²¹⁷ Just as a set of agents with a number of deontological commitments may be better at, say, maximizing the greatest happiness for the greatest number, than a group that worked explicitly with such a maxim.

standards that the other members will follow, but the same result is often produced more effectively by the shared belief that there is a language which all of the members are speaking, and a right or wrong way of speaking it. The attitude most people have towards the apparently arbitrary norms of their culture and time (styles of dress etc.) makes it clear that we are born with a tendency to think that we are not merely conforming to some arbitrary standard, but rather that there are ‘objective standards’ that we are collectively tracking. Our conception of language is the product of this objectifying response when exposed to linguistic behavior. This conception forms a background against which deference takes place, and seems, most of the time, entirely natural. It is a background in which there are such things such as the French language and the English language, and there are such things as speaking these languages correctly and incorrectly. These assumptions are frowned upon by linguists, often for good reason, but they play a role in the unreflective behavior of even those speakers who have been taught that these background assumptions are not justified. The concept of meaning is tied up in a conceptual web with this pretheoretical conception of language, and the associated conception of meanings as being generally shared by all speakers of a language -- which is responsible for our “slovenly” attributions of content -- is a result of this.²¹⁸

Deference is thus important not because shared concepts are somehow needed to provide a causal explanation of our deference behavior;²¹⁹ rather, a community’s having

²¹⁸ Issues relating to this are discussed, to a certain extent, in appendix F.

²¹⁹ Something like this understanding of the appeal to deference is suggested by Bilgrami in the following:

The argument from deference proposes that we explain it by adopting Burge’s social externalism about concepts and contents. The claim is that the fact that Bert (and all of us) so readily changes his ways to conform to the standards of his community and its experts is to be explained by saying that he all along had the expert’s concept. He was just misapplying it. He now has come to have a fuller understanding of it. (Bilgrami 1992a, 78-9.)

While Burge does point out deference as an important phenomenon, it is difficult to find in “Individualism and the Mental” anything like the claim that deference is a phenomenon that we need shared concepts to explain. A *causal* account of deference behavior could still be provided in individualistic terms. The ‘argument from deference’ is not that Bert’s changing his ways “to conform to the standards of his community and its experts is to be explained by saying that he all along had the expert’s concept” (Bilgrami 1992a, 78-9), but rather that Bert’s changing his ways to conform to the

the right sort of deference is partially constitutive²²⁰ of their having a shared system of concepts. The difference between the ‘right’ and the ‘wrong’ sort of deference can be understood in terms of speakers’ self-interpretations, and whether or not the imperative to defer is perceived as being hypothetical or categorical.²²¹

The attempts to ‘explain away’ deference and our belief ascriptions through appeals to pragmatic considerations seem unsuccessful, since they presuppose an indefensible picture of the interpretee’s psychology. As a result, our ascriptional practices seem to favor non-individualism over both individualism and anti-individualism.

1.23 Dismissing the practice

Finally, the individualist can argue that our everyday ascriptions do not carry significant weight when discussing philosophical/theoretical issues in the theory of meaning. A fairly clear example of this attitude can be found in Bilgrami’s *Belief and Meaning*:²²²

The appeal to ordinary practice of reporting seems hardly to acknowledge that the question here is a thoroughly philosophical one. The notion of content which is the subject of philosophical question is a notion which has a theoretical status because the very issue underlying the philosophical question is a theoretical one: what notion of content to posit such that it will serve in explanations of behavior as well as stand in essential relations with the external world? Why should we take facts about ordinary, everyday reports to be relevant to this theoretical notion? If Burge’s examples tell us anything, it is that ordinary, everyday practice is neglectful of the fairly serious differences between Bert and his fellows.²²³

This seems to deny that our ascriptions have even *prima facie* weight.²²⁴ However, the mere fact that one is looking for a *theory* to account for the truth of some of our semantic terms’ general characterizations does not give one the right to simply dismiss the *practice*

standards of his community is to be explained by his having beliefs (about doctors’ knowing more about arthritis than him, etc.) that also allow us to say that he has misapplied his concept if his usage does not correspond the community standard.

²²⁰ In Davidsonian terms, we could say that, their having a shared system of concepts is a matter of their being best interpreted as having such a shared system, and their deferring as they do (with the corresponding self-interpretation) contributes to just such an interpretation being the best.

²²¹ The fact that we do defer -- and see ourselves as responding to categorical rather than hypothetical imperatives -- plays a crucial part in a number of arguments against the individualist’s appeal to general beliefs in section 2.

²²² Once again, it should be noted that while Bilgrami could be understood as defending individualism here, his actual position is best understood as ascriptionally non-individualistic.

²²³ Bilgrami 1992a, 74. See also pp. 5-6.

²²⁴ “Ordinary chat ... is not in itself of relevance or interest in the philosophical study of meaning.” (Bilgrami 1992a, 75.)

that constitutes the putative extensions of those terms. Unless, of course, one is willing to simply change the subject (which some have accused Bilgrami of doing).²²⁵

Still, our actual practice of using our terms is not sacrosanct. For terms like “witch” or “phlogiston” it seems as if our past applications were entirely mistaken. Even though our ascriptions suggest that *other things being equal* a non-individualistic interpretation would be better, if other things are not equal (in particular, if philosophical/theoretical considerations favor the individualist), they should not carry any decisive weight. It is possible that enough deeply entrenched beliefs are incompatible with non-individualism that it could not be accepted into any equilibrium, and this possibility is the subject of the next section.

2. The general characterizations

The fact that deeply entrenched aspects of our ascriptional practices are non-individualistic does not entail that there will be a non-individualistic equilibrium available for our semantic terms. While the fact that we make non-individualistic ascriptions provides an at least *prima facie* case for thinking that such ascriptions are correct, the individualist can still argue that, even if such ascriptions make up part of the putative extension of our semantic terms, they could not be part of any final equilibrium between our ascriptional practices and the general characterizations we associate with our semantic terms. The individualist will insist that individualism cannot be given up once we see how closely it is linked to the general beliefs about content. The debate between individualism and anti-/non-individualism has often been understood as between those motivated by a desire to stick closely to our ascriptional practices (the putative extensions) and those motivated by more ‘theoretical’ concerns (the general characterizations).²²⁶ If it turns out that equally well entrenched general characterizations

²²⁵ See Travis 1995, 144.

²²⁶ Bilgrami, quoted above, is the most explicit about characterizing the debate this way, but the view is fairly widespread. Arguments for non-individualism often amount to little more than a statement of

are incompatible with non-individualism, one might not be able to find *any* equilibrium for our semantic terms, and some form of bifurcationism would seem necessary. There is, then, considerable reason to want to accommodate as many of the general beliefs associated with our semantic terms as possible within a non-individualist framework. This section will argue that many of the most central of these beliefs can, in fact, be so accommodated. Not only can many of the general beliefs that seem incompatible with non-individualism be understood in a way that is compatible with it, but the understandings of such beliefs that remain incompatible with non-individualism are also incompatible with the Kripke cases. Furthermore, there are a number of general beliefs that fit better with non-individualism than individualism.

2.1 We speak a shared language

One general belief that clearly seems to fit better with the non-individualist picture is the belief that we speak a shared language (ie: for most speakers we meet, we feel that we are speaking the same language as them -- though some may speak it differently). While such a belief is clearly quite deeply held, it is not at all obvious what it is. Just what is it to share a language in the first place? The individualist, with his commitment to each speaker having his own independent idiolect,²²⁷ cannot take language to be shared in any robust sense. As a result individualists tend to come up with an *ersatz* notion of a shared language, based on sufficient similarity, or simply reject the notion of a robustly shared language as an obviously confused idea. Indeed, the two responses tend to go together. Individualists claim that we understand ourselves as speaking a shared language in a way that is obviously false, and the *ersatz* notion is all that can be salvaged from the linguistic phenomena.

Burge's thought experiments, and those against usually depend upon its incompatibility with various deeply entrenched beliefs we have about self-knowledge, behavioral explanation etc.

²²⁷ Once again, the important part of his position relates not to our each having our own idiolect, but to the *idependence* of these idiolects.

The individualist will argue that while people tend to believe that (almost) everyone in the community speaks the same way, they greatly underestimate the differences between speakers (differences that can be, after all, considerable).²²⁸ In light of this, the individualist will recommend that we understand ‘sharing’ a language in a way that does not require that the languages involved be *exactly* the same (since facts about pronunciation etc. (not to mention malapropisms etc.) make it clear that no two people speak in the same way). Rather, a ‘language’ should be understood as ‘shared’ by a population if the idiolects of the speakers in the population are similar enough for them to communicate effectively, they learn how to speak from one another, are ‘corrected’ by each other, etc. Languages would thus be shared to a greater or lesser degree depending upon how much the idiolects overlap, with a limiting case of a language being fully shared by groups of speakers whose idiolects are indistinguishable. This limit would, of course, probably never be reached, so there would remain a sense in which languages are never fully shared.²²⁹

Nevertheless, even those people who are willing to recognize differences in pronunciation often tend to think that we at least *mean* the same thing by our words as our fellows, and the assumption that we refer to the same things by our words as our fellows is not as obviously false as the corresponding idea about pronunciation.

This may be to a large extent because, as the discussion of malapropisms above should have made clear, what a speaker means by his terms can depend upon the usage in his community even if the actual (phonetic) word he attaches to the concept is different from the one so attached in his community. If I mean *epithet* when I say “epitaph” (and thus the meaning of “epitaph” in my idiolect is determined by how “epithet” is used by

²²⁸ Linguists often can find informants who will insist that everyone in the community pronounces the words in the language the way they do, and when the linguist provides alternate pronunciations found among other informants, the informant in question will insist that such forms are incorrect and that no one speaks that way.

²²⁹ “If we are precise about what constitutes a language, it is probably the case that no two people actually do speak the same language.” (Davidson 1992, 260, see also Chomsky 1988 36-7, Lewis 1969, 201.)

my community), there is a clear sense in which I share a language with my fellow speakers. Even if we are not speaking the same way, there is still a “division of linguistic labor” in such cases. This “division of labor” metaphor suggests a way in which speakers can share a language without speaking the same way.²³⁰

Indeed, not only is speaking the same way not necessary for sharing a language, it is not sufficient either. Imagine, for instance, that two historically isolated communities lived in the same environment, but never interacted with each other (and were, perhaps, even unaware of each other’s existence). By sheer coincidence, the two communities came to speak in exactly the same way, applying the same words to the same things etc. Now, even if one were to admit that these speakers spoke *exactly* the same language,²³¹ there would remain a sense in which they did not *share* a language at all. One can understand languages as shared in the sense that I ‘share’ a haircut with everyone who has a haircut sufficiently similar to mine (a sense in which I probably ‘fully’ share my haircut with no one), but one can also understand languages as shared in the sense that a cooperative is shared by its members. The individual member has a stake in the cooperative, and the actions of the other members determine just what cooperative he has a share in in a way in which other people who ‘share’ my haircut do not affect the haircut I have. It is in this latter, more robust, conception of sharing that the two communities do not share a language at all, while members of the same community who speak differently can share a language in this sense.

In spite of its ‘robustness’ this conception of sharing a language is, if anything, more attainable than the individualist’s similarity-based notion according to which no two

²³⁰ The division of linguistic labor is something the ascriptional (rather than the methodological) individualist cannot really account for.

²³¹ And this admission is not one we should clearly make. One could compare the discussion in chapter one of the disagreement between a molecular and an evolutionary biologist over what counts as being of the ‘same species.’ Something molecularly identical to our dogs but having no evolutionary connection to them would count as a dog for the first but not for the second. In much the same way, a language that had the same phonology, syntax and even semantics as English, but which had a completely different origin, might count as being the same language as English to a formal linguist but a different language to a historical one.

people ‘fully’ share a language. I ‘share’ a language with other members of my community because I look to their help in deciding how to apply my own terms and they do the same with me. Because of this, once the idea of sharing a language is tied to our shared practice and less closely to the idea of speaking in exactly the same way, the idea that we share a language seems neither obviously false, nor accountable for in individualistic terms. We can understand people having the same language in terms of their sharing a language, or understand people sharing a language in terms of their speaking in the same way. The former strategy clearly seems to be the more workable.

While the non-individualist can claim that we really do *share* our language, since we rely upon others to specify the referents of our terms etc., the individualist’s understanding of the shared language belief requires that we understand it in a way that is obviously false (i.e.: we all speak in exactly the same way) and then tone it down so that it is roughly true (we all speak in roughly the same way). The individualistic understanding of the shared language belief requires, then, that we understand it as false. Since the belief is comparatively well entrenched, this should count against individualism.

2.2 Our terms have determinate extensions

Another general belief that may not sit well with individualism relates to the assumption that the extensions of our terms are more determinate than our discriminating capacities: we sometimes do not know whether a term applies to a certain object or not. This belief is phrased in metalinguistic terms, but it is manifested in non-metalinguistic attitudes such as, say, realizing that, while one is sure that a particular bird one is looking at (in good light etc.) is a raven or a crow, one does not know which.²³² Indeed, I initially wasn’t even sure if there was a difference between these two types of birds. I asked

²³² See also appendix F for a further discussion of how such attitudes could be manifested.

around, and while everyone was sure there that was a difference, no one was sure what it was (actually got conflicting suggestions about which one was bigger).

Such a belief causes problems for the individualist because there are many cases where the speaker will not be willing to say that a given term definitely does or does not apply to a given item. These may occasionally be colloquial way of reporting vagueness (“I’m not sure if Bill is bald or not”), but frequently they are not. If content is specified individualistically, and I am inclined to say of a particular bird, “I don’t know whether that bird is a crow or a raven,” then my term “crow” should be neither true nor false of that particular bird. If we determine content individualistically, many of our terms will be undefined for items we intuitively take to be either in or out of the term’s extension.

This fact is obscured if we focus on cases like Bert, who has a fairly specific, albeit idiosyncratic, conception of what arthritis is. Individualism would seem more plausible if we believed that people were more semantically decisive than they often are. The fact that people’s discriminatory capacities seem to underdetermine what we intuitively take them to be referring to may, ultimately, be a more serious problem for the individualist than the fact that they often misidentify what is intuitively taken to be the referent of their terms. Not only will the individualist leave us with thoughts of an intolerably general character, but he will have to treat us as simply confused when we claim not to know whether, say, a particular bird is a crow or a raven. Such claims presuppose that what a crow is extends beyond our dispositions to classify things as crows or not, and the individualist must say that, if I am not disposed to apply the term or its negation to the bird in question, I am mistaken in thinking that there is a fact of the matter as to whether or not a particular bird is a crow.

There may, of course, be cases where our uncertainty about whether to apply the term corresponds to an uncertainty about whether the term applies at all. We may be unsure whether, say, a habit is a type of dress, but we may be inclined to think that this kind of uncertainty corresponds to an indeterminacy or vagueness of the facts in question. Such

cases, seem, however, quite distinct from my not knowing whether a particular bird is a crow or an raven even though I'm quite sure that its one or the other. The former case is one where we may think that, even for someone who has fully mastered the term, whether it applies to a given item is not settled. The latter case is one where we are inclined to think that we have just not fully mastered the term. The individualist would have to assimilate all uncertainty to the former type, while the two types of case seem quite distinct.

This belief that the extensions of our terms are more determinate than our discriminating capacities should not be confused with (though it can lead to) the belief that our terms have completely determinate extensions. This stronger belief requires that every concept either does, or does not, apply to any possible object.²³³ This stronger claim has seemed to many to express a belief that should be just given up.²³⁴ It is

²³³ The following examples from Frege and Derrida are fairly representative of how widespread this train of thought can be in philosophical thought about language:

This involves the requirement as regards concepts, that, for any argument, they shall have a truth-value as their value; it shall be determinate, for any object, whether it falls under the concept or not. In other words: as regards concepts we have a requirement of sharp delimitation; if this were not satisfied it would be impossible to set forth logical laws about them. (Frege 1891, 148.)

What philosopher ever since there were philosophers, what logician ever since there were logicians, what theoretician ever renounced this axiom: in the order of concepts (for we are speaking of concepts and not of the colors of clouds or the taste of certain chewing gums), when a distinction cannot be rigorous or precise, it is not a distinction at all. (Derrida 1977, 123-4.)

²³⁴ It has been a target of Wittgenstein's (especially in the discussion of "family resemblances" in Wittgenstein 1953.), and Searle (1977), among others, has characterized it as simply a hangover from positivism. Furthermore, prototype-based accounts of our concepts (see Lakoff 1987, Rosch 1975) are based on precisely the assumption that the Frege/Derrida hypothesis is false and that our concepts do not have such determinate extensions. For instance, Jackendoff stresses that there is considerable evidence that our application of words is often governed not by our having a set of necessary and sufficient conditions for a term's application, but rather by our having a prototype associated with the word, with the result that we are comfortable applying the word to items that match the prototype closely and are increasingly unsure about the term's application the further items get from the prototype in question. (Jackendoff 1988, 1990.) While we certainly need not follow Jackendoff in abandoning the extensional approach to semantics, even within the extensional approach, the prototype phenomena could be seen as affecting the shape of the resulting theory. In effect, the prototype theory might suggest that most terms be treated as vague or 'fuzzy,' and that whatever extensional analysis one comes up with for terms like 'bald' or 'heap' should be extended to the other terms in the speaker's idiolect as well. The resulting theory would seem to predict and explain his behavior perfectly well. Nevertheless, Frege would be the last person to be impressed by the purported findings of prototype theories, and he might, with some justification, argue that psychological facts about how we are disposed to apply our terms no more determine what concepts are than psychological facts about how we are disposed to make inferences determine which inferences are, in fact, good ones. All

important, however, to recognize that the weaker claim, that our terms' extensions are more determinate than what our discriminatory capacities can pick out, is enough to cause problems for individualism; one does not need the stronger claim that all of our terms have completely determinate extensions. One need not commit oneself to the claim that we have no vague terms in the language to think that, even if I only know that gold is "an expensive yellow metal," I mean more by "gold" than "expensive yellow metal."

2.3 There is a fact about what we mean

One problem with tying what we mean to social usage is that it may seem to undermine the factuality of meaning. If what we meant were determined purely by facts about our dispositions or composition, then the facts about what we meant would be tied exclusively to things about which there seem to be settled facts. On the other hand, if meaning is tied to what our community means, then if there is no fact about, say, what community we belong to (what 'sociolect' we are speaking), then there will be no fact about what we mean.

This problem is heightened by the fact that anti-/non-individualistic positions are often associated with advocating 'the primacy of the sociolect.' Such accounts suggest that the norms that the individual's idiolect is responsible to are determined by 'the usage of the community,' 'accepted usage,' 'the public language,' 'the sociolect,' etc. If the speaker's idiolect represents his "always partial, and often in part incorrect, understanding of his language,"²³⁵ then his 'language' should be something that can be specified independently of his idiolect. Such accounts suggest that we should be able to give an account of just what these public languages are, and doing so is notoriously difficult. The non-individualistic picture suggests that there will be a single 'community' whose 'use' determines the sociolect that all of its members speak. But such 'languages,' while they are supported by our idea of, say, French and English being different

these facts could show is that our linguistic behavior seriously underdetermines which language we speak and which set of concepts we are, in fact, using.

²³⁵ Dummett 1975, 31.

languages, are notoriously hard to individuate. Davidson points out that “there is a problem of deciding what group is to determine the norms,”²³⁶ and Chomsky elaborates on a similar point:

In ordinary usage . . . when we speak of a language, we have in mind some kind of social phenomenon, a shared property of a community. What kind of community? There is no clear answer to this question. We speak of Chinese as a language, whereas Spanish, Catalan, Portuguese, Italian and other Romance languages are different languages. But the so-called dialects of Chinese are as varied as the Romance Languages. We call Dutch a language and German a different language, but the variety of German spoken near the Dutch border can be understood by speakers of Dutch who live nearby, though not by speakers of German in more remote areas. The term “language” as used in ordinary discourse involves obscure sociopolitical and normative factors. It is doubtful that we can give a coherent account of how the term is actually used.²³⁷

Not only are public languages hard to individuate, but the idea that each of us belongs to a group of people all of whom speak the same language is not as easy to justify as it might at first seem. The people I defer to may not, in fact, be the same people I recognize as speaking the same language as me. The idea that there may be a single self-standing community all and only the members of which speak the same language may have to be given up. Who an individual defers to may be different for any given term, and whether the relevant ‘experts’ will always be in the same community is doubtful. The experts I would ultimately defer to for the use of my scientific terms may be members of the international scientific community who do not even speak English, while the people I am willing to defer to for the use of other terms may not extend beyond my city or even my block.²³⁸ There may be no ‘community’ the use of which can be taken as authoritative for all my terms, and there may be no other person who recognizes just the same experts for all his terms as I do. It might seem, then, that there can be a separate practice associated with each of our words, and that a given speaker can be a part of the same practice as his neighbor with respect to some of his words while not with others.

²³⁶ Davidson 1991b, 197.

²³⁷ Chomsky 1988, 36-7.

²³⁸ Once again, Chomsky presents this worry in a particularly succinct fashion: “Suppose the expert happens to speak a variety of Cockney that I can barely understand? Is that part of my sociolect? I can perfectly well defer to my Italian gardener with whom I share only the Latin names for elms and beeches. Is his Naples dialect part of my sociolect?” (Chomsky, letter to Bilgrami, cited in Bilgrami 1992a, 266.)

This does not mean, however, that there are any serious problems for non-individualistic conceptions of content. Rather, it just shows that such conceptions of content should be divorced from a naive picture of what a ‘sociolect’ must be, or even from the idea that the ‘sociolect’ is an essential part of understanding the ‘social’ aspect of language. There may be a fact of the matter as to which group a speaker would ultimately defer to for each of his terms without there being a recognizable group to which the speaker belongs all of whom speak exactly the same ‘sociolect.’ Non-individualism and the ‘social character of language’ require only that we at least occasionally defer to, and hold ourselves responsible to, the usage of others. It does not require that those others be understandable as speaking the same ‘sociolect’ as we do.

Of course it may also be the case that there may be no fact of the matter as to who the speaker intends to defer to, and in such cases, the anti/non-individualistic picture would seem to suggest that there is no fact of the matter as to what we mean. However, in such cases, the claim that there is no fact of the matter as to what the speaker means has some plausibility. For instance, if a speaker is unaware of the fact that doctors and lawyers define “narcotic” differently and is equally inclined to defer to each group of experts, there may be no fact of the matter as to what the speaker means by the term, since the two expert communities define the term differently. Such cases, where a term is defined differently by different expert communities, need not be that uncommon. One could also cite physicists vs. chemists on what counts as a “molecule,” nutritionists vs. botanists on what counts as a vegetable, British and American English on what counts as an “endive,” etc.

2.4 The explanation of behavior

Perhaps one of the most deeply entrenched general characterizations related to the belief that our behavior is explained by, among other things, our beliefs and desires.²³⁹ While one may occasionally be able to deal with a conflict between a term's putative extension and general characterization by treating the relevant aspect of the general characterization as simply mistaken (just as, say, the belief that fish were a natural kind could be simply given up), it would be hard to take this line with the claim that our behavior is explained by our beliefs and desires. Most philosophers (the perhaps not most psychologists) who urge rejecting *this* belief would probably deny that there are beliefs and desires at all.²⁴⁰ As a result, if our non-individualistic ascriptional practices really were incompatible with our using belief ascriptions to explain our behavior, then we would have to either give up those practices, accept some sort of bifurcationism, or deny that we have any beliefs at all. Behavioral explanation, then, might provide the most serious obstacle to finding a non-individualistic equilibrium. The belief that our attitudes explain our behavior is extremely well entrenched, but it is also often taken to be incompatible with non-individualism.

One line of argument to the effect that non-individualistic ascriptions cannot be used to explain behavior runs as follows: (1) whatever (causally) explains our behavior must supervene on our internal states, and (2) if non-individualistic belief ascriptions were correct, then beliefs would not supervene upon our internal states, so (3) if non-individualistic ascriptions were correct, beliefs would not (causally) explain our behavior.²⁴¹ However, just because mental states are not individuated by their causal

²³⁹ This should not, of course, be understood as an empirical generalization of the type of explanation we actually give when we explain someone's behavior. Indeed, many, if not most, such explanations may make no explicit reference to beliefs or desires. "He's hungry," "She's drunk," "It's her day off," "His father just died," "She's a vegetarian" etc. are all perfectly good explanations of behavior even though they make no explicit reference to an agent's beliefs or desires. Rather, the claim is that our beliefs and desires at least *can* serve as elements in a (causal) explanation of our behavior.

²⁴⁰ For instance, some 'eliminative materialists' advocate denying that there really are beliefs and desires precisely because they take 'belief/desire theory' not to explain our behavior nearly as well as its hoped-for cognitive-psychological successors. (See Churchland 1981, Stich 1983.)

²⁴¹ Fodor, for instance, argues raises just such considerations in the following passage:

powers, it does not follow that they cannot figure in causal explanations. Chairs are not, after all, individuated by their causal powers, and this does not prevent there being causal explanations involving chairs. Just because ‘relational facts’ determine what counts as a “chair” it doesn’t follow that there cannot be causal explanations involving such terms. After all, it is a perfectly good causal explanation to say that John died because he was run over by a Mercedes Benz, even if being a Mercedes Benz does not supervene on the physical properties of the car that hit him.

Another reason for thinking that individuating content non-individualistically provides belief contents poorly suited for the purposes of behavior explanation, a reason that does not presuppose a particular picture of causal explanation, run roughly as follows. If the contents of an agent’s thoughts are not tied to his conception of what he purports to be thinking about, and his behavior is explained by his conception of things, then those thought contents will not explain his behavior. Bert’s behavior, for instance, might seem better explained by attributing to him beliefs involving *tharthritis* rather than *arthritis*, because *tharthritis* more closely matches his conception of what “arthritis” is. To take another example from Burge, if an agent has a conception of what a “contract” is that involves the belief that contracts must be in writing, then attributing to such an agent a desire to have a *contract* with his neighbor will not explain his insistence that he and his neighbor write down their agreement, though attributing to him a desire for a *thontract* (just like a contract but must be written) will.

These considerations are, however, far from convincing. First of all, the examples presented above might suggest that there will always be a set of things corresponding to the agent’s conception (the set of all aches in the joints and limbs for *tharthritis*, and the set of all written contracts for *thontract*). However, this is not always the case. Not only

If you are interested in causal explanation, it would be mad to distinguish between Oscar’s *mental* states and Oscar2’s; their mental states have identical causal powers. But common sense deploys a taxonomy that *does* distinguish between the mental states of Oscar and Oscar2. So the commonsense taxonomy won’t do for the purposes of psychology. Q.E.D. (Fodor 1987, 34.)

is the agent's conception often too vague to pick out any determinate set of objects (see 2.2), but there is often *nothing* that corresponds to such a conception. For instance, in addition to believing that other people have beliefs about "contracts," the "contract" believer may also believe that, say, there are classes in contract law taught at law schools, etc. Since there are no classes taught in *thcontract* law, there will, in fact, be nothing that corresponds to his "contract" conception. If any set of items is going to be picked out for his term, part of his conception will have to be understood as mistaken, and the non-individualist will argue that the belief that contracts must be written is precisely the belief that should be given up. So the non-individualist can argue that his attributions can actually do a *better* job at capturing the agent's conception of things than the individualist's.

Furthermore, the purported problems with behavioral explanation are already present once we accept the ascriptional practices that Kripke appealed to,²⁴² and the problem that there may, in fact, be nothing that corresponds to the agent's conception is most evident in the case of proper names. For instance, if I believe (falsely) that Bert is a doctor, my calling Bert if I break a finger might seem to be better explained by a belief about *Thert* (i.e.: someone who fits my conception of Bert), but that still doesn't change the fact that I have a false belief about *Bert* and not a true belief about some 'intentional entity' who happens to fit my conception of Bert.²⁴³ Whatever problems non-individualistic assignments produce for behavioral explanation are already there when we assign truth-conditions in accordance with non-negotiable intuitions about proper-name reference.

Nevertheless, the fact that the problems with behavioral explanation are also caused by accepting Kripke's work on proper names does not show that they are not real problems. And if the purported problems could not be dealt with, then the non-negotiable Kripke-intuitions and the equally non-negotiable belief that our behavior can be

²⁴² Much the same could be said of the Putnam practices (our 'water' behavior being purportedly better explained by attributing a concept that picks out both H₂O and XYZ).

²⁴³ Indeed, even if there were someone else who 'matched' my conception of Bert, I would still be referring to Bert.

explained by our beliefs and desires would force us at least to split up belief content. As discussed in chapter one, this would involve arguing that there are, in effect, two types of content: a “narrow content,” that is individuated internally (individualistically), and a “wide content,” that is individuated externally (non-individualistically).²⁴⁴ The former is meant to capture the agent’s perspective on things and thus explains his behavior; the latter captures our perspective on what the agent says and thus corresponds to what we evaluate as true or false.²⁴⁵

Since we are trying to find a single equilibrium, this is clearly not the kind of solution we are looking for. However, while such a position clearly isn’t an ideal one, some have argued that it couldn’t even be a *possibly* acceptable. Bilgrami has criticized such solutions for sacrificing what he calls the “unity of content,” namely, the idea that it is the very same contents that are externally determined and evaluated as true and false that serve in the explanation of behavior.²⁴⁶ And his ‘unity thesis’ can be understood as an insistence that some single equilibrium be found, if any are to be found at all. If the unity thesis really is non-negotiable, then the prospect of finding *separate* equilibria for capturing truth conditions and point of view seems fairly grim.²⁴⁷

Given that this is the case, it may be worth taking a closer look at whether non-individualistic content attributions really don’t allow us to explain behavior. In particular, one might hope that there is a way to capture the agent’s conception of things while still individuating his concepts non-individualistically. There may be a way to *reconcile* non-individualism and behavioral explanation, without *isolating* them from

²⁴⁴ Various forms of which are found in, for instance, Putnam 1981a, Block 1986, McGinn 1982, Loar 1985, and Field 1978.

²⁴⁵ Such accounts presuppose, of course, a particularly restrictive account of how the behavior to be explained can be characterized. In particular, narrow contents can serve to explain what an agent is *trying* to do, but they seem poorly suited to explain mistaken behavior or what an agent succeeds in doing (correctly calling a tree an “elm”, etc.). (For a discussion of this, see Brandom 1994.)

²⁴⁶ “What makes it possible for intentional states to have a role in the commonsense psychological explanation of behavior is understood as criticizing the idea that there could be a separate equilibrium capturing our ‘point of view’ which was isolated from concerns relating to what this point of view was a point of view on.

each other the way the bifurcationist does.²⁴⁸ Perhaps the most obvious way to do this is to make a more extensive appeal to the agent's beliefs. While the above-mentioned speaker's insistence that he get his neighbor to sign a contract in writing may be explainable by attributing to him a desire for a *thontract*, the behavior seems to be explained equally well by attributing to him *both* the desire for a contract *and* the belief that contract must be in writing. My phoning Bert when I have a broken finger is explained not by a desire to call *Thert*, but rather by a desire to phone Bert and a belief that Bert is a doctor. Non-individualistically individuated beliefs can still explain behavior as long as they are taken together with other beliefs.²⁴⁹

Bilgrami, however, claims that such accounts, like those that rely on a distinction between narrow and wide content, are bifurcationist and thus violate his unity constraint.²⁵⁰ However, while such accounts do separate the roles of meaning and belief (the first giving truth conditions and the second explaining behavior), it does not follow that Bilgrami's criticism is justified. While the unity thesis seems to match at least our pre-theoretical intuitions about the role of belief content in behavioral explanation,²⁵¹ it only requires that the externally determined contents be the "*same* contents which go into the commonsense psychological explanation or rationalization of behavior."²⁵² This requirement has both a weak and a strong reading. The weak reading is that the contents

²⁴⁸ Reconciliation would show how non-individualistically specified content can be used to explain behavior, while isolation would rely on dividing the explanatory work types of content do in such a fashion that non-individualistically specified contents need play no role in behavioral explanations.

²⁴⁹ Something like this position seems to be taken by, say, Davidson and McDowell. Davidson's actual position on this issue is not entirely clear and on questions of Davidson exegesis, I will be following Bilgrami at this point (see Bilgrami 1992a, 177-80 for a discussion of remarks made by Davidson in conversations relating to this issue). Furthermore, while Davidson uses this strategy to reconcile behavioral explanation with non-social externalism, he does not use it to defend social externalism, which, of course, he doesn't endorse.

²⁵⁰ "[For McDowell], there are *concepts* and *conceptions* obtained from looking at the truth theories and beliefs respectively. How can one deny that this amounts to bifurcationism, with one notion of content containing the concept and the other, the conceptions?" (Bilgrami 1992a, 172.)

²⁵¹ Though Stich suggests (in the last chapter of Stich 1983) that it may turn out to be false, and that the 'beliefs' that are part of our verbal reports are not in the same system as those that explain our behavior. If this turned out to be true, and one did not, like Stich, think that this meant that there were no beliefs, one would have a type of explanatory bifurcationism in which the contents of both components were externally determined.

²⁵² Bilgrami 1992a, 3.

that are externally determined serve as *part* of an explanation of the agent's behavior. The strong reading is that the externally determined theory of meaning for an agent should provide *all* the resources needed to explain the agent's behavior: if you know what he means by a word, you know all his associated beliefs²⁵³ as well, and thus can explain all his behavior.²⁵⁴ It should be clear that those who propose dividing up the explanatory work between narrow and wide content violate the unity thesis in both its weak and its strong forms, and are thus bifurcationists in both senses. However, while someone who splits the explanatory work between belief and meaning may violate the strong version of the unity thesis (since he maintains that a theory of belief is also needed to explain behavior) he need not violate it in its weaker form. It *is* the externally determined contents that go into the explanation of behavior; they just do so by providing the content of the beliefs in the belief component. The conceptions associated with the belief theory are in no sense independent of the contents associated with the truth theory, and thus there is not the type of independence associated with theories that distinguish narrow and wide content. So the question becomes, then, which form of the unity thesis should we accept? Given that both the strong and the weak versions of the unity thesis are able to account for the intuition that "what makes it possible for intentional states to have a role in commonsense psychological explanation of behavior ... is the very fact of their possessing content which is externally constituted,"²⁵⁵ the weaker version seems preferable. While the goal of making a person intelligible is a generally admirable constraint on a theory of interpretation, it is not clear that one should want to apply it directly to the theory of meaning. After all, it is quite unintuitive to claim that simply knowing what an agent means by his terms is enough to understand, and possibly even

²⁵³ Or perhaps all the beliefs relevant to a particular 'explanatory locality.' (Bilgrami 1992a.)

²⁵⁴ Though, of course, the thesis is often (but not in Bilgrami's case) accompanied by the corollary that we never 'really' know what anyone means by their words and that 'complete' communication is impossible.

²⁵⁵ Bilgrami 1992a, 15.

predict, his behavior.²⁵⁶ It is much more natural to think that we can predict or explain an agent's behavior because we have some idea of what he believes, and this would correspond to our having a theory of belief rather than a theory of meaning. Of course, by constructing the type of theory of meaning that is in line with the strong unity thesis, one would also come to know all the agent's beliefs as well. So the issue is not whether we need to know the agent's beliefs, but rather whether or not these beliefs should be kept theoretically distinct (but not isolated) from the meaning theory.

This brings us to a crucial difference between theories that divide the explanatory work between narrow and wide content and those that divide it between meaning and belief: while the former posit a split within single folk semantic/psychological categories such as meaning and content, the latter mirror the commonsense conception that we mean certain things by our words and we have certain beliefs about the world and that while these two phenomena are related, they are distinct.²⁵⁷ As a result, while strongly unified theories have an advantage over bifurcationist ones in that they can give a unified account of pre-theoretically unified subject matter like meaning, they have a disadvantage compared to weakly unified accounts in that they also unify subject matters that are pre-theoretically distinct.²⁵⁸

However, it may seem that pointing out that we can distinguish what a speaker means from what he believes may not, in itself, be enough to allow non-individualistically individuated beliefs to explain behavior adequately. Having the content of an agent's

²⁵⁶ The idea that the content of a single belief should be enough to explain a certain piece of behavior seems undermotivated. Admittedly, single beliefs and desires are cited to explain behavior (my opening an umbrella is explained by my belief that it is about to rain etc.), but this explanation usually works against a potential background of other beliefs and desires (that I believe umbrellas stop water, that I don't want to get wet etc.) that must be presupposed for the explanation to work. We don't have any problem assuming such a background, so there seems to be no reason to want to incorporate all this background information into the content of the belief itself, especially since such an incorporation leads to a number of serious problems that will not, however, be discussed here (see my "Comments on Bilgrami's *Belief and Meaning*" MS 1993).

²⁵⁷ This need not, of course, involve a commitment to any sharp analytic/synthetic distinction.

²⁵⁸ Of course, theoretical progress may necessitate redrawing some of the boundaries on our pre-theoretical conceptual map, and it is always possible to claim that the dictates of theory determine how the conceptual map is to be drawn. Nevertheless, if Bilgrami is to take this line against the weak unity theorist, he will undermine his own case against the bifurcationist.

beliefs determined by things other than his conceptions would seem to allow us to attribute beliefs to him that do not seem to make sense of his behavior. For instance, if Bert really believes that he has arthritis, rather than *tharthritis*, then he believes that he has a rheumatoid ailment that affects the joints exclusively, but attributing *this* belief to him will make no sense of his insistence that there is nothing wrong with his joints. Appealing to his belief that one can get ‘arthritis’ in the thigh does not help if this can be understood as the belief that you can get a rheumatoid ailment exclusively affecting the joints in the thigh. The mere fact that we can appeal to other beliefs, then, does not necessarily show that there is nothing unproblematic with non-individualistic belief attributions.

This is, once again, a phenomenon we also see with proper names. If Jocasta is Oedipus’s mother, and Oedipus intends to marry Jocasta, then Oedipus intends to marry his mother. Yet, since Oedipus also intends to do whatever he can to avoid committing incest, his eagerness to enter into such a marriage seems mysterious. If we are to make sense of his behavior, then, it may seem that we must deny that Oedipus intends to marry his mother, and doing this may seem to return us to attributing only beliefs that are in line with the agent’s conception of the objects of his beliefs.²⁵⁹ However, while many explanations of behavior require that such a constraint be satisfied, behavior is occasionally explained in a way that explicitly violates the constraint. For instance, we can explain why Oedipus married his mother by pointing out that he believed both that, since she was queen of Thebes, his mother would make a great match, and that his mother was not a relative of his (and certainly not his mother). Oedipus himself might reject such characterizations (claiming instead that he believed that, since she was queen of Thebes, *Jocasta* would make a great match, and that *Jocasta* was not a relative of his),

²⁵⁹ Bilgrami formulates such a requirement in his “constraint C”: “When fixing an externally determined concept of an agent, one must do so by looking to indexically formulated utterances of the agent which express indexical contents containing that concept and then *picking that external determinant for the concept which is in consonance with the other contents which have been fixed for the agent.*” (Bilgrami 1992a, 5.)

but there is a perfectly good sense in which such ascriptions explain his behavior (which consisted, after all, in his marrying *his mother*).

Ascriptions of the sort just discussed are, of course, commonly referred to as *de re*, and are more ‘perspicuously’ put on the following lines: “Oedipus believed of *his mother* that *she* would make a great match” and “Oedipus believed of *his mother* that *she* was not a relative of his.” Such *de re* ascriptions are clearly not constrained by the speaker’s conception of things, and the ascriptions that are so constrained are typically referred to as *de dicto*.²⁶⁰ Both types of ascription are used to explain behavior (the *de dicto* ascriptions explaining what he is trying to do (marrying Jocasta), the *de re* ascriptions explaining what he does (marrying his mother)), and just as we can say that Oedipus believes of his mother that she would make a great wife, we can say that Bert believes of *a rheumatoid ailment exclusively affecting the joints* that he has *it* in his thigh. If we take an agent’s conception of things to constrain the type of *de dicto* ascriptions we can make of him rather than constraining the contents of the beliefs themselves, then one can individuate beliefs non-individualistically while still allowing that there are no true *de dicto* belief-ascriptions that do not match his conceptions.

It may seem that relying on the *de dicto/de re* distinction in this way also produces a type of bifurcationism that violates the unity constraint: the beliefs as ascribed *de dicto* explain behavior while those ascribed *de re* are evaluated as true or false. However, not only can we evaluate beliefs as ascribed *de dicto* as true or false, and use the beliefs we ascribe *de re* to explain behavior, but the two types of ascriptions should also not be understood as corresponding to two types of belief. If the ‘bifurcationism’ is *perspectival* rather than *substantive*, it isn’t clear that the unity constraint is violated. Remember, the unity constraint only requires that the externally determined contents are the “*same* contents which go into the commonsense psychological explanation or rationalization of

²⁶⁰ I.e., if the *ascriber* would accept the characterization of the belief, it is *de dicto*, if the *ascriber*, but not the *ascriber*, accepts the characterization, it is *de re*. (See Dennett 1982, Brandom 1994.) For a fuller discussion of this, see appendix E.

behavior,”²⁶¹ and the fact that these contents do fill their various roles when viewed from different perspectives doesn’t change the fact that it is the same contents viewed in each case.²⁶²

It seems, then, that there is no compelling reason to think that non-individualistically individuated thought contents are unable to serve in behavioral explanations, and thus no reason to think that this issue should prejudice us towards believing either that there could be no equilibrium points or that any equilibrium points would have to be individualistic.

2.5 Self-knowledge

Non-individualistic content ascriptions are also often taken to be incompatible with the deeply entrenched general belief that we are authoritative about the content of our beliefs and utterances. Non-individualistic content attributions are supposed to be incompatible with our having knowledge of the contents of our own thoughts, since they allow that the content of our beliefs can be determined by facts of which we are unaware. If what we mean by our words and the content of our thoughts are determined by contingent factors that we may know nothing about (e.g.: what the atomic structure of gold really is, how “arthritis” is used in the community), then we often don’t seem to be in a position to know what we mean. As Bilgrami puts it:

The idea is that a pair of internally identical twins on earth and twin earth respectively have different but, as far as *they* can tell, indistinguishable substances, which they both call “water,” in their environments; so, given externalism, they have different ‘water’-concepts and ‘water’-thoughts. Since they have different thoughts without really being able to tell the difference, they do not fully know what their own thoughts are.²⁶³

The claim that we know what we mean may seem easier to dismiss than the claim that our beliefs explain our behavior.²⁶⁴ The idea that we have such ‘privileged access’ to

²⁶¹ Bilgrami 1992a, 3.

²⁶² Once again, the *de dicto/de re* distinction is discussed in greater detail in appendix E.

²⁶³ Bilgrami 1992b, 235.

²⁶⁴ Woodfield, for instance, takes it as relatively unproblematic to say that “a third person might well be in a better position than the subject to know which object the subject is thinking about, hence be better placed in that respect to know which thought it was.” (Woodfield 1982b, viii.)

the contents of our thoughts has occasionally been dismissed as just a relic of our Cartesian heritage that, like much of that heritage, we should rid ourselves of as quickly as possible.²⁶⁵ While such a gambit might have some plausibility to someone who cannot tolerate any position that even *reminds* him of Cartesianism, it does not seem, ultimately, to be very convincing. One wants a picture in which we have self-knowledge in *some* sense, and so finding an interpretation of self-knowledge that does not involve an appeal to Cartesian privileged access is preferable to using the rejection of such access as an excuse for giving up self-knowledge. There has been no shortage of papers generated by this topic,²⁶⁶ but perhaps the most prominent defender of the compatibility of non-individualism and self-knowledge is Tyler Burge. Burge argues that, since the very same factors that determine what we mean also determine what we take ourselves to mean, the fact that we may not be authoritative about these factors does not prevent us from being authoritative about what we mean.

Burge argues for their compatibility in the following fashion:

Knowledge of one's own mental events ... consists in a reflexive judgment which involves thinking a first-order thought that the judgment itself is about. The reflexive judgment simply inherits the content of the first order thought. Consider the thought, 'I hereby judge that water is a liquid.' What one needs in order to think this thought knowledgeably is to be able to think the first-order, empirical thought (that water is a liquid) and to ascribe it to oneself, simultaneously. Knowing one's thoughts no more requires separate investigation of the conditions that make the judgment possible than knowing what one perceives.²⁶⁷

Such 'homophonic' explanations of self-knowledge seem in many ways unconvincing. All such arguments tell us is that our second order judgments about our thoughts and assertions about what we mean will be both true and warranted (since what determines

²⁶⁵ There has also been, of course, considerable evidence from psychology that suggests that we are not nearly as authoritative about our thoughts as we might have thought.

²⁶⁶ See, for instance, Falvey and Owens 1994 and the papers cited therein.

²⁶⁷ Burge 1988, 656. Davidson makes essentially the same point in Davidson 1989, 168, and Davidson 1994b, 235. However, given that Davidson endorses this sort of argument, it isn't clear why he thinks that there is "a conflict between Burge's social externalism, which ties a speaker's meaning to an élite usage he may not be aware of, and first person authority" (Davidson 1991b, 197), since the argument form he uses above works equally well for social externalism.²⁶⁸ For a discussion of similar problems relating to attempts to explain our knowledge of language as knowledge of an "M-theory", see Dummett 1975. It should also be noted that, as Joe Camp has stressed, even if the homophonic account were able to account for our positive self-knowledge, we would still lack semantic self-knowledge in the *negative* sense. We could not be certain that any of our claims not to believe a particular proposition would be true.

the contents of these second order judgments and assertions is manifestly what determines the contents of the first order thoughts and utterances they are about). However, the truth and warrant of such second-order judgments and assertions does not amount to self-knowledge about our first-order thoughts and assertions unless we know what we are thinking and saying when we make our second-order judgments and assertions. Someone who is inclined to think that non-individualism threatens our knowledge at the first level is unlikely to admit that such knowledge of our second-level thoughts and utterances is unproblematic. The homophonic argument merely seems to push the problem back a level.²⁶⁸

That the homophonic argument may not preserve all of the self-knowledge we intuitively take ourselves to have can be seen in the fact that it still leaves open the possibility that we can't be certain that we are thinking at all. If we are thinking, our self-ascriptions will have access to the same contents as our thoughts, but if we simply seem to be thinking, then our self-ascriptions are just so much empty noise.²⁶⁹ Still, this type of skeptical worry about lack of self-knowledge is not peculiar to non-individualistic positions; rather, it is the product of any 'externalistic' account.

Perhaps a more serious worry about the compatibility of non-individualism and self-knowledge that also does not seem to be addressed by the homophonic defense of self-knowledge given above is that non-individualistic ascriptions seem to require that the interpretee have thoughts that, were he both rational and aware of their content, he would not have. As a result, if one wants to hold on to the idea that the interpretee is rational, one must understand him as not being fully aware of the content of his thoughts. If Bert really believes that he has a disease of only the joints in his thigh, then either he is

²⁶⁹ This is, of course, why the anti-skeptical arguments commonly attributed to Putnam and Davidson are unacceptable as 'straight' solutions to skepticism about our knowledge of the external world. (The attribution is considerably more plausible in Davidson's case than Putnam's.) If one is a Cartesian, one can be sure that one is thinking, but within Putnam and Davidson's framework, one has no such guarantee. (This is abundantly clear in Davidson 1987 443-44, and Putnam 1981c, 17.) The extent to which semantic externalism can be incorporated into an argument against skepticism is discussed further in my imaginatively titled "Externalism and Skepticism" (MS 1996).

irrational or he doesn't know what he is thinking. Since we have no reason to think that he is irrational, we must assume that he doesn't know what he is thinking.²⁷⁰

However, it should be clear that we would have the same problem with ascriptions such as "Oedipus believes that his mother is not his mother." Furthermore, it should also be clear that, if *de dicto* and *de re* ascriptions are distinguished the way they were in the previous section,²⁷¹ there will be no problem with non-individualistic ascriptions resulting in our ascribing to the interpretee blatantly inconsistent *de dicto* beliefs. The constraint on *de dicto* attributions (that they involve characterizations of the objects of belief that the interpretee accepts) will ensure that, provided the interpretee is rational, the beliefs attributed will not be obviously inconsistent.²⁷² Our 'self-knowledge' of our beliefs as ascribed *de dicto* becomes, from this perspective, rather trivial, since it is built into the structure of the *de dicto* ascription process that we would be able to recognize and accept the ascribed beliefs as our own. We would not, of course, have any special knowledge or 'first person authority' about our beliefs as they are ascribed *de re*, but this follows directly from the account as well.²⁷³

²⁷⁰ This line of thought is found most clearly in Bilgrami. See, for instance 1992b, 240.

²⁷¹ And appendix E.

²⁷² One can thus make ascriptions non-individualistically and still accept most of the following claim of Bilgrami's:

What is it that you are attributing when you say that he believes he has arthritis in his thigh, what does 'arthritis' in the *de dicto* specification of the content tell us? ... The prohibition of substitutions in *de dicto* attributions turn on the attributee *not knowing* of the co-extensiveness of the expressions involved; and I am precisely protesting that there must be some, at the minimum one, and often more, expressions that can be substituted, or else it is not clear that *any* concept of arthritis should be attributed to him. And moreover if we admit that 'a disease of the joints only' is not one of the expressions to be substituted for 'arthritis' because the attributee does not know of the co-extensiveness, then ... that puts into doubt the relevance of Putnam's externalist concept of arthritis to *de dicto* attributions. (Bilgrami 1992b, 245-6, see also, Bilgrami 1992a, 43.)

It is only the last sentence that the non-individualist need have problems with. If *de dicto* ascriptions corresponded to a special kind of belief, then showing that *de dicto* attributions must be constrained by the agent's 'point of view' would seem to show that there was a type of belief content that was so constrained. However, if *de dicto* ascriptions are only a type of *ascription* for the *very same* beliefs that can be ascribed *de re*, then constraints on these *de dicto* ascriptions would not entail constraints on the beliefs themselves. Because of this, while his ignorance that one can't get arthritis in one's thigh prevents one from substituting "a disease of the joints" into *de dicto* ascriptions of 'arthritis' beliefs to Bert, the substitution is allowed when we make *de re* ascriptions of those same beliefs, and thus a far less 'constrained' externalism is permitted in characterizing the contents of his thoughts.

²⁷³ One might, of course, argue that this avoids the problem only by introducing another type of bifurcationism, but, once again, it is not at all clear whether a *perspectival* bifurcationism of this sort is

Still, there is *some* sense in which the non-individualist can allow that we remain authoritative about even our *de re* beliefs. In particular, while Bert may not know that arthritis affects only the joints, the fact that he means *arthritis* by “arthritis” is still determined by other attitudes of his -- his believing doctors know more about arthritis than he does, his willingness to understand himself as mistaken about arthritis when he is corrected, etc. If he lacked such attitudes, there would be little reason to claim that he had made a mistake about *arthritis*. Unlike the anti-individualist (who will see Bert as mistaken no matter what his other attitudes are), the non-individualist leaves what determines whether Bert is to be interpreted in accordance with communal usage or not to be determined, ultimately, by facts about *Bert*. In this sense, Bert is still authoritative about the content of his beliefs even as specified *de re*.

Finally, it should be pointed out that ascriptively individualistic accounts of content are themselves surprisingly unfriendly to self knowledge. After all, as discussed in 1.221 and 1.222, our attributions of content to ourselves are guided by just the same principles as the attributions of content we typically make when interpreting others. Bert is more likely to say that he used to believe falsely that he had arthritis in his thigh than he is to say that he had a true belief to the effect that he had *tharthritis* in his thigh and a false belief that he meant what others did by the term. Because of this, while the individualist may avoid saying that speakers are ignorant of some aspect of the content of their beliefs, he seems forced to say that they are often (especially in cases relating to deference or the acquisition of new information) mistaken about the content of their beliefs and about what beliefs they have had in the past. If Bert did not have *any* beliefs about arthritis until he was corrected by his doctor, all of his claims about past beliefs involving ‘arthritis’ will turn out to be *false*. While mistakes about what one meant in the more distant past can occur, cases like Bert’s being corrected would involve him being

at all incompatible with the unity thesis. Especially since it corresponds with a *de re/de dicto* distinction that is motivated independently of trying to reconcile non-individualism and self-knowledge.

mistaken about what he meant right up until the time of the actual correction. Indeed, the fact that he will defer (and, crucially, defer categorically) when corrected suggests that he didn't even know what he meant at the time of utterance. If we did defer for merely pragmatic reasons and did operate under hypothetical imperatives to speak as others do, then individualistic content attributions would not pose a threat to self-knowledge. It is only because deference cannot always be explained the way Davidson and Bilgrami propose (see 1.222 above) that self-knowledge becomes a problem for the individualist. Had we always deferred for the sorts of reasons that Bilgrami suggest, there wouldn't be a problem about our not knowing what we meant in such cases for the individualist, since our self-ascriptions would be correct by the individualist's lights.

Indeed, there is good reason to think that individualism is ultimately *more* harmful to our conception of self-knowledge than non-individualism. On the non-individualistic account, Bert's claiming that he has arthritis in his thigh does not need to be explained by his not knowing what he meant by "arthritis" -- it can simply be explained by his having a false belief, namely, that one could get arthritis in one's thigh. It is not, however, clear that the individualist can claim that Bert's mistaken assertion that he used to mean *arthritis* (rather than *tharthritis*) by "arthritis" can be explained by anything other than his having a mistaken belief about what he meant; there seems to be no available alternative belief about the world that could be used to explain his behavior.²⁷⁴

This brings us back to something that came up in section 1.23. Bilgrami was unimpressed by appeals to our typical attributions of beliefs and desires, since he believes that when working on the *theory* of meaning, such rough-and-ready attributions should not necessarily be taken seriously.²⁷⁵ However, the requirement that a theory must

²⁷⁴ Bilgrami seems right to say that it is "foolish and unintuitive" to "grant that a wholly *non-psychological* phenomenon, i.e., an entirely abstract philosophical doctrine about reference, should be seen as raising an obstruction to self-knowledge" (Bilgrami 1992b, 236), but it may be the case that it is his own theory, and not non-individualism, that provides such an obstruction.

²⁷⁵ Such an attitude is connected to Bilgrami's commitment to putting the 'theory' back into the theory of meaning, a commitment evident in passages such as the following.

attribute beliefs that are compatible with self-knowledge undermines the possibility of taking such an unashamedly theoretical stand. Our ‘everyday understanding’ of others will be very much like our everyday understanding of ourselves, and so any theory that takes our everyday understanding of others as irrelevant runs the risk of providing a theory according to which we systematically misunderstand ourselves, and thus lack self-knowledge. Self-knowledge thus provides a constraint that prevents the appeal to the theoretical status of the theory of meaning from being applied in too *ad hoc* a fashion.

It seems, then, that considerations of self-knowledge, while often presented as one of the stronger considerations in favor of individualism, may actually provide us with a compelling reason to reject it. As a result, considerations of self knowledge present no real obstacle to our finding an equilibrium point incorporating the ascriptional practices Burge appeals to.

3. Conclusion

It has been argued in this chapter that while the incorporation of the socially infected ascriptions that Burge appeals to into an equilibrium for our semantic terms is considerably more controversial than the incorporation of the ascriptions appealed to by Kripke and Putnam, such an incorporation remains preferable to trying to find an equilibrium that does not endorse the truth of such ascriptions. It was argued that such ascriptional practices are not only themselves fairly deeply entrenched but are also not clearly isolatable from the extremely well entrenched practices that Kripke appeals to. Attempts to dismiss or ‘explain away’ non-individualistic ascriptions by appealing to more ‘pragmatic’ factors were shown to rely upon a number of unrealistic assumptions about what we take ourselves to be doing when we make ascriptions or defer to

I do not have in mind everyday understanding of each other but, rather, what underlies everyday understanding or what everyday understanding consists in. The question, what it is for agents to have the concepts and contents which we attribute to them in our everyday understanding of them, cannot be answered by looking to the way in which we attribute those concepts and contents to them in everyday situations. (Bilgrami 1992a, 5-6.)

correction, and even those ascriptional practices that might have seemed to favor individualism (those associated with slips, malapropisms and misunderstandings) were shown to be best understood within an non-individualistic framework.

General beliefs associated with our speaking a shared language whose terms have determinate extensions favor incorporating the Burge phenomena, and it has been argued that the most frequently cited candidates for entrenched general beliefs incompatible with non-individualistic ascriptions (those relating to self-knowledge and behavioral explanation) turn out to be compatible with them after all. Indeed, accounting for self-knowledge ultimately turns out to cause more problems for the individualist. Reconciling our ascriptional practices with these general beliefs has involved picking up on some of the themes introduced in chapter two by stressing (1) that the practices are *non-individualistic* rather than *anti-individualistic*, (2) that the *de dicto/de re* distinction is a distinction between types of ascriptions, not between types of beliefs, (3) that the position is *methodologically*, rather than *ascriptionally*, individualistic, and (4) that the position relies heavily on our *self*-interpretations to fund the non-individualistic ascriptions. (Considerations similar to these four will turn up again when we discuss the possibility of incorporating the Wilson ascriptions in the next chapter.) There seems good reason to think that such strategies could also be used to reconcile the practices with any other seriously entrenched general belief. As a result, while the prospects for finding an ascriptionally individualistic equilibrium seem unpromising, there does not seem to be any reason to think that we cannot find an ascriptionally non-individualistic equilibrium for our semantic terms.