

1. Putative extensions and general characterizations

I have called many things “cats,” my parents had three, my neighbor has two, I see them constantly in pet stores, on the street, etc. I also have a number of beliefs about cats: they have a certain appearance, they are animals, they are part of a biological kind closely related to lions and tigers, they like milk, they don’t like getting wet, they can climb trees and fall great distances, etc. This is not unusual. For most¹ terms in my language there is a range of cases in which I would be willing to apply them, and a number of (at least implicit) beliefs I have about the type of thing the terms are supposed to pick out.² I’m not unusual in this respect. Each of us has a number of beliefs associated with a given term and is disposed to apply it in a range of cases (though these beliefs and applications may vary from person to person). We can call the conditions that prompt the use of a given term its “putative extension,”³ and those more general beliefs about what one is talking about the “general characterization” associated with that term.⁴

The general characterizations will thus include all of the sentences we would endorse that include the word in question. They are, very roughly, our beliefs about or pertaining to the types of things that fall within the putative extension. For instance, when we refer to things as “chairs,” we may have a number of beliefs about them (that they are pieces of furniture, people sit on them, they have backs, they have legs, etc.), and these characterizations are not independent of the putative extension. Indeed, they are often a product of our encounters with (and occasional reflections upon) the things that fall within the putative extension. Furthermore, while one can include all of the agent’s beliefs in the

¹ I leave open the possibility that a different story may need to be told for connectives, etc.

² For other discussions of these two elements in a words use, see Brandom 1994 (especially pp. 119, 129, 427), Dummett 1973, and Sellars 1954.

³ We may want to include something like the term’s perceived anti-extension under the general heading of the putative extension (i.e.: not just the things we are sure are x’s but also those things we are sure are not x’s).

⁴ It should also be noted that the general characterizations may just be presuppositions rather than explicitly held beliefs, and they may also be fairly undefined. (See appendix F.)

general characterization, there will be some privileged sub-set of these characterizations that will be considered less negotiable than the others.

For instance, for an “observational” term such as, say, “pen”, one discovers a term’s putative extension by discovering the stimuli the presence of which prompt the term’s application. One discovers the term’s general characterization by discovering what ‘inferential transitions’ its user is disposed to make with it. The former is manifested in a speaker’s ability to pick out, say, a pen when asked to fetch one. The latter is manifested in his ability to describe pens, his belief that one can write with pens, that ball-points and rollerballs are both pens, etc. Both of these sets of abilities are important parts of a speaker’s linguistic competence, and while the two usually go together,⁵ they seem to be psychologically distinct.⁶

The putative extension of an observational term like “pen” might, then, bear some resemblance to Quine’s notion of “stimulus meaning,” which Quine defines as our “total battery of present dispositions to be prompted to assent or to dissent from the sentence.”⁷ Quine believes that the stimulus meaning of a given term can do more or less justice to its ‘real’ meaning. For an observation term like “red” the stimulus meaning may do “full justice” to its meaning, while for a term like “bachelor” (for that the speaker’s use relies heavily on “collateral information”)⁸ it does not. In much the same way, we might say that the putative extension of a given term can be more or less independent of the general beliefs. In the case of theoretical terms, the putative extension may be tied quite closely to, and even dependent upon, aspects of the general characterization, and this may also be true of the more ‘inferential’ applications of observational terms: there are situations where I

⁵ But not always; one could learn to identify pens very well without knowing what purpose they served, or know quite a bit about their function and history without necessarily being able to recognize them.

⁶ For instance, one can have one set of abilities impaired without severely damaging the other. (See Marconi 1995 (esp. pp. 141-3) for descriptions of cases where speakers have lost much of their ability to identify objects while their ability to associate names with descriptions or synonyms remained unimpaired.)

⁷ Quine 1960, 39. Though, technically, the putative extensions bear more of a resemblance to Davidson’s modification of stimulus meaning, with the ‘stimulus’ being distal rather than proximal. (See Davidson 1990.)

⁸ Quine 1960, 42.

may say a cat is present just because of general beliefs I have about the noises they make, the types of tracks they leave, etc.

It should be noted that putative extensions and general characterizations each have semantic programs associated with them. “Inferential-role semantics” focuses on the general characterizations associated with a term, while “information-based semantics” focuses on the term’s putative extension. That is to say, these programs, while perhaps not *identifying* meaning with one of the two aspects of our linguistic practice, still treats it as *determined by* the relevant aspect (for a discussion of this, see section 3.2).

2. Tension between the two

There is thus a range of cases in which we apply a given word and certain general beliefs we have about the type of thing the word applies to. These two come into conflict when we discover that there are items to which we apply a given term but that do not correspond to our general conception of the type of thing that term should apply to.⁹ In such cases we can either stop applying the term to the object in question or we can alter our general conception of the type of thing the word applies to. An instance of the first alternative can be seen in the development of the use of “fish.” We have a set of general beliefs about fish such as: they live in the water, they have gills rather than lungs, are cold blooded, have scales, don’t suckle their young, etc. The fact that whales don’t satisfy the last four of these descriptions caused at least some of us to change the putative extension of “fish.”¹⁰ On the other hand, if the general characterizations are too far out of line with the putative extension, they will usually be revised or given up. For instance, we would probably give up the claim that chairs have legs rather than claim that all of the modern legless chairs are not ‘really’ chairs. In some cases, the entire framework of the *type* of

⁹ Or vice versa. It should also be noted that there may be conflicts within these two factors themselves. We may apply our terms in an irregular manner (sometimes applying the term to an object, sometimes not) and our general ideas about the type of thing the term applies to may be inconsistent.

¹⁰ This description is perhaps a little misleading, since it is given at the social level, and the change in usage of an entire society will take longer and not proceed in such a simple fashion.

thing one is talking about is thrown out when it turns out to be in conflict with too much of the putative extension.

A fairly clear instance of this is our eventually giving up the assumptions that trees or lilies form any type of 'natural kind.' Deciduous trees and evergreens bear no close taxonomic relation to each other, and many lilies are more closely related to onions than they are to other lillies, but we don't say that evergreen's aren't 'really' trees or that onions are really a type of lilly. The same applies to "reptile" as well (crocodiles being cladistic sisters of birds, not lizards). Of course it is worth noting that one can take an extremely hard line, and keep with the natural kind intuition in these cases and just dump the whole practice ("[The] cladogram of "higher" primates contains an interesting implication: there is no such thing as an ape, at least as usually defined").¹¹

Once such a basic belief is given up, the type of belief the general characterizations must focus on will change. On perhaps an even more basic level, changes in what we take to be involved in being, say, a natural kind can reverberate throughout the entire system. For instance, if we change our conception of what a biological kind is (possibly because of our previously having dealt with an isolated range of cases) we may have to reconsider the extensions of many other terms.

We may say that an aspect of our usage is "entrenched" to the extent that it is difficult to dislodge when it comes into conflict with other aspects of the term's usage, and which element of a speaker's usage is given up is a function of the comparative entrenchment of these elements. Aspects of our usage are frequently entrenched simply by being long-standing, but many other factors can entrench a belief or application. The extent to which it is supported by or supports other entrenched beliefs (including one's theoretical or ideological commitments) could, for instance, contribute to the comparative entrenchment of some aspect of a term's use.

¹¹ Gould 1983, 358.

Still, entrenchment is, to a large extent, a function of numerical support. Fifty beliefs or applications will, collectively, tend to be more deeply entrenched than any single belief or application. For the most part, if a small percentage of the cases that fall within the putative extension conflict with the general characterization, then those cases will tend to be removed from the putative extension, while if a large percentage of them do, it will be the generalized characterization that is modified. Of course, deciding what to give up will not always be a simple matter of finding out what percentage of the items in question satisfy a particular general characterization. Some items that a term is applied to (and some aspects of its characterization) may be more deeply entrenched than others. For instance, some of the examples of what I call “chair” are more central (paradigmatic/prototypical) than others. If our general characterization of chairs ruled out the bean-bag chairs in my parents’ basement, I might be willing to say that beanbags were not ‘really’ chairs. On the other hand, if the characterization ruled out the formal chairs in my parents’ dining room, I would probably say that there was something wrong with the characterization.¹² In much the same way, some of the general beliefs about chairs may seem more central than others: that chairs must have legs is something that can be (and probably has been) given up, while that we can sit in chairs may be less negotiable.¹³ The comparative entrenchment of various elements within a term’s putative extension and general characterization may vary from person to person, and it may be, at least to a certain extent, indeterminate. Which elements we will be willing to give up may depend upon the order with which we are faced with various conflicts. (The significance of this will be discussed in chapter four.) When there are no conflicting elements left in a terms usage, that term’s usage can be said to be at an “equilibrium,” and a terms ‘real’ extension can be identified with the putative extension it would have at such an equilibrium. (This will be discussed in greater detail below.)

¹² And this would not change even if bean bags were, in fact, more common than formal chairs.

¹³ Of course, it may turn out that there are no strictly necessary and sufficient conditions of being a chair, in which case one would look for some larger set of conditions most of which would have to be satisfied for something to be a chair. (See Rosche 1975, Jackendoff 1988.)

Once this picture is in place, it should be clear that there will be a number of intermediate cases where it is not clear which should be modified, the general characterization or the putative extension. Indeed, the whale/fish case is one where different speakers can have different reactions as to which element should be modified.

For example, Linnaeus¹⁴ may have initially put whales within the putative extension of “fish.” However, he also believed that fish and mammals were mutually exclusive natural kinds. On investigating what distinguished these kinds, he discovered that, while mammals had lungs and warm blood, fish had gills and cold blood. These beliefs about what distinguished fish from mammals were incompatible with the inclusion of whales within the putative extension of “fish,” and since these beliefs were more deeply entrenched than the application of the term to whales, Linnaeus came to understand that application as mistaken. He came to the realization that whales were not (and never had been) fish.

On the other hand, if the disposition to call whales “fish,” or the belief that “fish” picks out a *functional* kind, is more deeply entrenched than the belief that “fish” and “mammal” pick out mutually exclusive natural kinds, one will characterize the latter belief as mistaken. Such a stance is taken by Melville’s Ishmael:

In his *System of Nature*, AD. 1766, Linnaeus declares, ‘I hereby separate the whales from the fish.’ But of my own knowledge, I know that down to the year 1850, sharks and shad, alewives and herring, against Linnaeus’s express edict, were still found dividing the possession of the same seas with the Leviathan. The grounds upon which Linnaeus would fain to have banished the whales from the waters, he states as follows: ‘On account of their warm bilocular heart, their lungs, their movable eyelids, their hollow ears, penem intransentem feminam mammis lactantem’ and finally ‘ex lege naturae jure meritoque’ I submitted all this to my friends Simeon Macey and Charley Coffin, of Nantucket, both messmates of mine in a certain voyage, and they united in the opinion that the reasons set forth were altogether insufficient. Charley profanely hinted they were humbug.... I take the good old fashioned ground that the whale is a fish, and call upon holy Jonah to back me. This fundamental thing settled, the next point is, in what internal respect does the whale differ from other fish. Above, Linnaeus has given you those items. But in brief, they are these: lungs and warm blood; whereas, all other fish are lungless and cold blooded.¹⁵

Different interests on the part of various speakers will make different aspects of the general characterizations and putative extensions seem more central than others, fishermen and

¹⁴ Swedish botanist (1707-78). This reconstruction of his thought process has no pretensions to historical accuracy.

¹⁵ Melville, 1983, 936.

scientists being interested in ‘functional’ and ‘biological’ classification respectively.¹⁶ These differences in what they hold to be central will lead fishermen and scientists to accept different ways of resolving conflicts between aspects of the term’s use. If the parties involved have such ‘incommensurable’ commitments, there may be no way to settle disputes about whether, say, it is really correct to call whales “fish.”¹⁷ In such cases, it may be better to say that the conflicting parties simply refer to different things by the term.

Fairly notorious instances of unresolved tension between the general characterization and the putative extension can be seen in cases such as our use of philosophically controversial terms such as “knowledge,” where some accept general characterizations that entail that we don’t really know anything (or know next to nothing), while others see this conclusion as grounds for reevaluating the general characterization and replacing it with a new one that is more in keeping with the putative extension.¹⁸ In these more deeply contested cases, it can seem particularly difficult to come up with a definitive answer as to how the tension is to be resolved. Certainly, one will not expect there to be anything like an algorithm that would tell one which of the two competing claims will have to be given up.¹⁹

Still, this will not always be the case in such disputes. In particular, if there is no equilibrium available corresponding to one party’s use, that party is probably best understood as simply mistaken. As a result, one may argue that, say, Ptolemaic astronomers meant what we do by “planet,” “star” etc. because, while they had a number

¹⁶ Mapping one’s terms on to natural kinds and mapping one’s terms on to functional kinds are both perfectly intelligible and frequently implementable projects, and one might have initially thought that one could do both with the same terms.

¹⁷ Indeed, even if we were, as Kripke and Putnam suggest, initially committed to treating terms like “fish” as natural kind terms, it has, of course, turned out that fish are not, on at least some accounts, a biological kind at all (lungfish being cladistic sisters with land animals, not the other fish). One should also note that Ishmael appeals to the testimony of friends from *Nantucket* (then the preeminent whaling town), since the appeal makes it clear that he is not simply being idiosyncratic, but is rather aligning himself with a different expert community (which includes, of course, “holy Jonah”).

¹⁸ Such terms may stand for what are, in W. B. Gallie’s terminology, “essentially contested concepts.” (See Gallie 1956.)

¹⁹ Nor should one expect that such issues to be settled by appealing to some causal fact that would settle what was really referred to (for a discussion of such theories, see appendix A).

of differently entrenched beliefs, there is no equilibrium governing those beliefs that differs from the equilibrium governing ours.

3. How both are central to meaning

While both the putative extensions and general characterizations play an important part in determining the meaning of our words, philosophical accounts of meaning often overemphasize one of the two. In particular both contemporary “inferential-role” and “information-based” semantics are unable to achieve a robust conception of objectivity because they focus upon one of the two factors and miss out on the dynamic relation between them.

3.1 Normativity

An excessive focus on either a term’s putative extension or its general characterization inevitably deprives our use of language of aspects of the normativity it (manifestly) seems to have. If the meaning of a term, “X,” is determined solely by its general characterization,²⁰ then it seems impossible for us to be wrong about the types of things Xs are. If we come to believe that, say, all reptiles have forked tongues, then that belief simply becomes part of what we mean by “reptile” and hence (analytically) true (though, of course, it could turn out that there are no reptiles).²¹ As a result of this, it could turn out

²⁰ A fairly clear instance of the overemphasis on the general characterizations can be seen in Mill’s theory of general names:

All concrete general names are connotative. The word *man*, for example, denotes Peter, Jane, John, and an indefinite number of other individuals, of whom, taken as a class, it is the name. But it is applied to them, because they possess, and to signify that they possess, certain attributes. These seem to be, corporeity, animal life, rationality, and a certain external form, which for distinction we call the human. Every existing thing which possessed all these attributes, would be called a man; and any thing which possessed none of them, or only one, or two, or even three of them without the fourth, would not be so called. (Mill 1843, 35.)

On such an account, the meanings of our general terms are determined exclusively by the general characterizations. If we happen to talk about actual items in the world, it is because they happen to fit the general characterization. Of course Russell’s theory of descriptions was eventually used to extend this connotative picture to proper names as well.

²¹ One can occasionally find philosophers willing to bite this bullet as well. Katz, for instance, claims that “if Putnam’s story about the things we have been calling cats turning out to be robots were true, it

that *none* of our applications of a given term are correct -- the connection between language and the world becomes more tenuous than it would intuitively seem. On the other hand, if the meaning of a term is determined solely by its putative extension,²² then it becomes increasingly difficult to see how we could ever misapply it.²³ Apparently substantive claims like “this is an X” must be understood as meaning roughly “this is a member of the set of things I call ‘X.’” The latter being a much less interesting claim that can be made true simply by its sincere utterance. Compare the differences between “She is a witch” and “Now she is what I’d call a witch.” The second, when sincerely uttered, might seem always to be true, but if meaning is determined solely by conditions of application, the truth-conditions of the first utterance must be understood in terms of the truth-conditions of the second.

Consider, for instance, the following two sentences:

(1) This is a cat.

(2) Cats are carnivorous.

If the meaning of “cat” is understood in terms of its general characterization, it is clear how a substantial claim can be made with (1), but a sincere utterance of (2) will merely be true in virtue of the meaning of “cat” (making it difficult to see how a substantial claim is made with it).²⁴ On the other hand, if the meaning of “cat” is understood in terms of its putative extension, then it is clear how a substantial claim can be made with (2), but sincere utterances of (1) acquire the same ‘analytic’ character that utterances of (2) did when

would follow simply that there never were any cats! There couldn’t be, since to be a cat something must be an animal.” (Katz 1975, 96.) (See also Lewis 1929, 26-7.)

²² Indeed, Quine’s notion of stimulus meaning does precisely this. Quine claims that the stimulus meaning of a sentence for a subject “sums up his dispositions to assent to or dissent from the sentence in response to present stimulation,” and that stimulus synonymy “can be used as a standard of translation not only for observation sentences but for occasion sentences generally” (Quine 1960, 34, 217). While Davidson differs from Quine in focusing on the distal rather than proximal stimuli, he shares with Quine a tendency to identify what our words *refer* to with whatever we are disposed to apply those words to.

²³ This has become very apparent in discussion of the “problem of error” (or “the disjunction problem” (Fodor), “the misrepresentation problem” (Dretske) or the “normativity problem” (Millikan)) in the works of ‘information semanticists’ who take truth-conditions solely to reflect the putative extensions of the terms in question.

²⁴ Of course this conclusion can be moderated a bit with ‘cluster’ description theories such as Searle’s, though the majority are still viewed as ‘collectively’ analytic. (See section 3.21.)

meaning was understood in terms of its general characterization.²⁵ General characterizations and putative extensions are each good at explaining how the utterances that characterize the other are contentful, so any account that focuses on just one of the two will not be able to account for the contentfulness of the very utterances it takes to be constitutive of meaning. The favored class of statements thus have the status conventionalists had assigned to the axioms of geometry. Viewed from outside the framework, they collectively defined the framework itself, but from within the framework, they express claims that are necessarily true in virtue of the constituted meanings. There is a sense in which tying meaning to general characterizations allows one to explain how we manage to say something with our utterances, but prevents one from adequately explaining how what we say is about anything, while tying meaning to putative extensions allows one to explain how our utterances are about objects in the world, but prevents one from explaining how we manage to say anything about them.

Of course, few people have suggested that general characterizations and putative extensions determine the meaning of *all* our terms. There is, however, a popular strategy of splitting the terms in the language into two groups, the first consisting of terms whose meaning is determined by their putative extension, the second consisting of terms whose meaning is determined by their general characterization. Bertrand Russell could be understood as presenting a theory with roughly this form. Logically proper names have their meaning determined by their putative extensions while descriptions, disguised or not, have their meaning determined by the associated general characterizations.²⁶ This leaves

²⁵ Some have tried to make this failing into a virtue and parlayed an account's inability to account for objectivity in this area into a positive anti-skeptical virtue of the account. (For a discussion of this, see appendix B.)

²⁶ Indeed, something like this structure is found in many if not most theories with a sharp distinction between 'observation' and 'theoretical' terms. This two-tiered structure can also be attached to various causal theories, with names, natural kind terms etc. having their meanings determined by the putative extensions associated with a term's baptism and with the other terms having their meaning determined by general characterizations the components of which eventually reduce to terms of the information-bearing variety.

We think that the description theory of reference fixing may well apply to 'pediatrician'; that it may well apply to non-basic "artefactual kind" terms; and that it may well apply to many other terms; for example, 'bachelor,' 'vixen' and 'hunter.' This is no threat to our theory of language, because the

the various types of ‘analyticity’ in place for different parts of the vocabulary, so while statements of type (1) could be made substantively with some terms, and statements of type (2) made substantively with others, there would be no terms for which statements of *both* type (1) and (2) could be made substantively. However, since “cat” seems to be used substantively in both (1) and (2), two-tiered theories do not capture our intuitions in this area.

A different type of ‘hybrid’ position equates a term’s meaning with a *combination* of its putative extension and general characterization (the first picking out the ‘wide’ and the second the ‘narrow’ content). However, it will be argued that simply *equating* wide content with putative extensions would seem to get the truth-conditions of our terms wrong (see 3.22), but if narrow content is determined by the general characterizations, it would seem that wide content would have to be determined by the putative extensions.

3.2 Attempts to salvage objectivity from within a single component

While there are obvious problems when one simply *equates* a term’s meaning with either its putative extension or its general characterization, it may seem that one could treat the term’s meaning as *determined* by one of the two and still capture our intuitive notion of objectivity. By having meaning determined by, rather than equated with, the chosen meaning-constitutive aspect, one leaves room for the idea that elements within the meaning-constitutive aspect could still count as being mistaken.

3.21 General characterizations and inferential-role semantics

That accounts based upon general characterizations had problems accounting for objectivity was not lost on everyone within this tradition, and some (e.g. Block) simply gave up on the idea that the type of ‘meaning’ that was determined by our general

reference of all the terms appealed to in a description theory must ultimately be explained causally. (Devitt & Sterelny 1987, 78.)

Of course, it isn’t at all clear that we really could find descriptions involving only natural kind terms which would correspond to every term in our language, and the suggestion that typical users of the language have the reference of their non-natural kind terms fixed this way seems incredible.

characterizations should, or could, concern itself with issues relating to truth. A less pessimistic line would be to try to find some notion of objectivity from *within* the general characterizations. This is the line taken by Searle when he moves from ‘straight’ to ‘cluster’ description theories in passages such as the following:

Suppose we ask the users of the name ‘Aristotle’ to state what they regard as certain essential and established facts about him. Their answers would be a set of uniquely referring descriptive statements. Now what I am arguing is that the descriptive force of ‘This is Aristotle’ is to assert that a sufficient but so far unspecified number of these statements are true of this object.... suppose we were to drop ‘Aristotle’ and use, say, ‘the teacher of Alexander,’ then it is a necessary truth that the man referred to is Alexander’s teacher -- but it is a contingent fact that Aristotle ever went into pedagogy (though I am suggesting it is a necessary fact that Aristotle has the logical sum, inclusive disjunction, of properties commonly attributed to him: any individual not having at least some of these properties could not be Aristotle).²⁷

What a term refers to can be understood in terms of what most of the items in the general characterization are true of, so any single element in the general characterization can be false, even if their disjunction is analytically true.

Of course, such attempts to find a notion of objectivity from within the general characterization have serious limitations, and Kripke has argued that, for instance, *all* of the general characterizations associated with a name could turn out to be false of its referent.²⁸

3.22 Putative extensions and information-based semantics

Dissatisfaction with accounts of meaning and reference that focus excessively on the general characterizations became more and more evident from the early seventies on, since it became generally accepted that our specifications could often underdetermine or even misidentify what we intuitively took ourselves to be talking about. If anything, accounts of meaning and reference now pay excessive attention to the putative extensions.

However, as discussed above, simply *identifying* what a term refers to with its putative extension has the serious drawback of being unable to find any substantial notion of objectivity with respect to obviously contingent utterances of ours such as “That is a cat.” How then is any normativity to be salvaged from the dispositions that determine a term’s

²⁷ Searle 1958, 94-5.

²⁸ See, for instance, his discussion of “Aristotle” in Kripke 1972.

putative extension? One popular strategy is to find some privileged set of dispositions that can be taken to be determinative of meaning, leaving open the possibility that the dispositions to apply our terms that are outside of the set will turn out to be misapplications. The two most popular candidates are those dispositions that accord with the dispositions of the community and those dispositions that we would have under certain ‘optimal’ conditions.²⁹ Neither of these, however, seems to adequately capture the normativity in question.

Consider first the ‘communitarian’ accounts that take the relevant dispositions to be those that agree with the dispositions of the community.³⁰ Such accounts require that, at the community level, a term’s putative extension just *is* its actual extension, and such a requirement seems quite unintuitive. Entire communities can be, indeed often have been, mistaken about the proper extension of particular terms in their language, so the communitarian accounts cannot fund the sort of full-blooded conception of objectivity we want. Furthermore, even if one were willing to accept that the community’s usage was incorrigible in this way (which is, of course, is an admission most of us are unwilling to make), the communitarian accounts only provide an explanation of an individual’s *non-perceptual* errors. Bert can be understood as mistaken in claiming that he has arthritis in his thigh because his disposition to apply “arthritis” to pains in his thigh does not match the dispositions of the community. Perceptual errors are less well accounted for. If Bert is disposed to apply the term “deer” to cows seen at a distance with antlers strapped to their heads, there is no reason to think that the communitarians will be able to account for this sort of systematic perceptual error. If cows with horns strapped to their heads look a lot like deer from a distance, then it will be true that other members of Bert’s community will be disposed to ‘misapply’ the term in the same way that he does.³¹ (Of course, the

²⁹ Boghossian 1989, 534.

³⁰ Such a position has been attributed to both Burge 1979 and Kripke 1982, though both attributions are questionable.

³¹ Boghossian 1989, 536.

community could tell that there was an error if it were allowed to ‘take back’ previous assertions when closer, but this type of story could be told about Bert on his own.)

The other alternative is to take the privileged set of dispositions to be those we would have under certain ‘optimal’ conditions (the conditions under which we learned the term, the conditions under which our cognitive mechanisms are functioning as they should, etc.).³² Such accounts have frequently focused upon accounting for putative ‘misrepresentations’ such as a frog’s ‘mistaking’ a BB for a fly,³³ but it seems unlikely that any such account will be able to explain errors such as Linnaeus's applying the term “fish” to whales. In particular, the putative extension of “fish” for (pre-1776) Linnaeus may be no different from the putative extension of “fish” for Ishmael. The fact that one correctly applies the term to whales while the other is making a mistake can, it seems, only be explained in terms of differences between their general characterizations. As a result, attempts to account for normativity by supplementing putative extensions exclusively with extra-conceptual factors will be unable to account for the difference between the two. Furthermore, the extensive literature on this sort of proposal shows a fairly clear consensus that the proposals on the table so far are all inadequate for solving the ‘disjunction problem.’ That is to say, they all fail to provide an explanation of how a dispositional account of meaning can leave room for the person with the dispositions ever making a mistake.³⁴

³² The cognitive system’s “functioning as it should” is here usually given some sort of ‘teleological’ explanation (ultimately to be explained by evolutionary biology). It should be noted, however, that one could understand a system as “functioning as it should” when all of its terms are at an equilibrium, though such a proposal would, of course, make essential use of the general characterizations.

³³ The frog’s much discussed response to small black objects in its environment may be a paradigmatically punctate disposition, and if human misrepresentation is to be understood on analogy with such cases, then general characterizations will play no essential role in accounting for it.

³⁴ It might be thought that such a strategy could be supplemented by some type of ‘metaphysical essentialism’ to get at a more full-blown notion of objectivity. An agent’s dispositions will pick out a certain putative extension, and the world’s ‘essential structure’ will determine which items are part of the term’s extension (i.e., those items that have the same essential structure as most of the privileged sample). Such accounts are discussed in greater detail in appendix A.

The dispositions that determine the putative extensions of our terms, like those that determine the general characterizations, cannot fund a sufficiently substantial notion of objectivity on their own, and so there is reason to look for another type of solution.

3.3 The need for a dynamic conception

Since we would like to see ourselves as making *substantive* claims *about* the world around us, a picture of meaning as the product of the *dynamic* between putative extensions and general characterizations seems desirable. *Saying* something *about* the world requires us to have both the putative extensions and the general characterizations in play for each word. Having each put checks on the other allows us to have a fairly robust conception of responsibility both to what it is for something to be an X (since aspects of a term's general characterization may be out of line with the majority of the items to which the term is applied), and to what items are, in fact, X's (since some of the items that the term is applied to may be out of line with deeply entrenched aspects of its general characterization). Any account of meaning in which a term's extension is determined by just one of the pair in question will be one in which there are no checks on the criteria in question, and will thus provide a comparatively anemic account of semantic normativity.

There may be other types of normativity that could be appealed to. Someone whose general characterization of a given word was satisfied by nothing could be viewed as violating the more 'pragmatic' norm that we should use words that actually pick out items in the world. In much the same way, someone whose putative extension for "gold" included iron pyrite could be seen as violating the pragmatic norm of having one's terms map on to natural kinds. In both cases the norms are pragmatic norms suggesting that we have made a 'mistake' by associating a particular meaning with a particular term, not semantic norms suggesting that we have made a mistake (a false belief or misapplication) *because* of what we mean by our term.

4. Equilibrium as a regulative ideal governing our linguistic practices

Our use of language is thus best understood as developing in a way that can be called, for want of a better word, “Socratic.”³⁵ This practice is regulated by the idea that, over time, conflicts between a term’s putative extension and general characterization will be resolved, and that one will end up with an ‘equilibrium’ in which the general characterization matches those items within the putative extension.³⁶ In some sense, meaning itself could be understood as corresponding to what our linguistic use would be in such an ‘ideal’ state. The idea that we do, in fact, mean something by our terms regulates our practice of trying to eliminate the inconsistencies between our ascriptional practices and our general beliefs. Of course the move towards equilibrium could be driven in part simply by our desires to keep our beliefs consistent. If we believe that all X’s are P, that *a* is an X, and that *a* is not P, then we will want to give up one of the three beliefs (the first of which could be part of our general characterization of X’s, while the second could be part of “X”’s putative extension). Nevertheless, our response to these conflicts is structured not *just* by a desire for consistency. While consistency can always be produced by *either* modifying the general characterization or altering the putative extension, in most instances we see one of these two as what is required given what we mean by the term. Which we choose is determined, of course, by how well entrenched the incompatible factors are within our practice. Since we view the process as getting at what we ‘really’ mean, not any way of eliminating the inconsistency will do.

Our assumption that we mean anything at all involves something like the regulative use of an idea of reason. There may not necessarily be anything corresponding to it, but it guides our investigations towards an independently desirable equilibrium. We do not see

³⁵ One can see instances of this ‘Socratic’ method in Hegel’s discussion of knowledge (Hegel 1807), Goodman’s discussion of induction (Goodman 1954), Rawls’ discussion of justice (Rawls 1971), and Burge’s account of linguistic development in general (Burge 1986c). Something like this methodology is suggested for semantic terms in Devitt 1994, 1996, though Devitt seems to build into his account a number of unjustified presuppositions about the methodology (Devitt’s account and some problems with it are discussed in appendix C).

³⁶ It should be stressed that this ideal regulates the use of individual speakers as well.

ourselves as simply being disposed to apply words to things that may or may not fit our beliefs about them, rather we understand ourselves as speaking a (shared)³⁷ language in which there are matters of fact about which general beliefs are true and false and which items we apply our terms to genuinely fall within their extension.

4.1 Conflicting pictures of linguistic development

Since we are committed to making our general characterizations and putative extensions cohere together whenever we find that they are out of step with one another, it can be tempting to think that there is a rationally coherent structure behind our practice, a pre-existent and independent structure against which our practice is evaluated. The regulative idea of meaning thus encourages a comparatively simple picture of linguistic development. In this picture there are roughly three things involved in a term's development: what the term actually means (what it refers to, its *actual* extension), what we take the term to apply to (its *putative* extension), and what we believe about what the term stands for. The simple picture sees linguistic development as arising when we realize that the putative extension of a given word doesn't match its actual extension and adjust our usage accordingly. For instance, when I was younger, I would have fairly comfortably classified mushrooms as a type of vegetable. When asked what vegetables I didn't like, mushrooms would have been near the top of my list, etc. I used to call mushrooms "vegetables," but I was eventually informed that mushrooms weren't vegetables, so I stopped calling them "vegetables." It can be tempting to see this part of my linguistic development as my simply bringing the putative extension of "vegetable" into line with the actual extension. According to this picture, linguistic development is merely a matter of chipping away at various misapplications of our words until we discover their 'true' extensions. Much the same can be said of the beliefs: we have a number of present beliefs, and we alter them so that they

³⁷ That we understand this language as shared is fairly important in explaining how we coordinate our language use with that of others. The suggestion that the idea of a shared language functions as a regulative ideal controlling our linguistic behavior is not only more phenomenologically plausible than more voluntaristic explanations such as Lewis's account of convention, but is also compatible with non-individualistic theories of content in a way in which conventionalist accounts are not. (For a more extended discussion of these issues see my "Convention and Language.")

correspond to what is actually true of the referent of the term in question. Such a picture is perhaps most compelling with natural kind terms, but it lies behind our thought about other tracts of our language as well.

We run into problems if we come to assume that there is something ‘out there’ corresponding to this regulative ideal, and the natural thought that both the conditions of application and the general characterizations are subject to correction can lead to a quasi-platonistic conception of language in which linguistic norms are robustly independent from our actual linguistic practice. (Indeed, something like this platonizing picture can be seen as motivating many, but by no means all, complaints to the effect that our language use is ‘degenerating’ or is being ‘corrupted.’) This quasi-platonist idea that meanings inhabit some practice-independent realm encourages the idea that they are available to us through some special ‘intuition’ and can thus be determined through introspection or some sort of ‘conceptual analysis.’

Such hopes are destined to be disappointed, and the fact that we cannot find anything in experience that corresponds to this regulative ideal can encourage another picture of linguistic development that goes to the opposite extreme. The rejection of the quasi-platonistic picture of linguistic norms, in which meanings are ‘out there’ and our practice is responsible to them, can leave one with a type of skepticism about the possibility of any sort of linguistic normativity that criticizes the practice itself. (I.e. it assumes that to criticize our practice we must be going beyond the practice itself.)³⁸ What we mean should be determined by our linguistic practice, so it may seem, at times, hard to understand what could constitute the ‘actual’ extensions that stand over-and-against the putative ones (which seem to reflect our practice), and if the independently determined extensions can’t be made sense of, they drop out of the picture, leaving the putative extensions on their own. What a speaker means by a term is understood, on such accounts, in terms of how he is disposed

³⁸ One way to correct the practice of the individual is to appeal to the practice of the surrounding community, but not only does this leave the objectivity problem at the communal level, but such accounts can seem like the arbitrary imposition of social usage upon the individual.

to apply it. Any attempt to get a more disposition-transcendent notion of meaning can thus come to be viewed as a type of unmotivated supernaturalism.³⁹

Such an account simply denies that linguistic development occurs at all. There may be linguistic *change*, but that is all. For instance, when I stopped calling mushrooms “vegetables,” I just changed what I meant by “vegetable.” While this picture has been defended even for natural kind terms, it is at its most plausible for non-natural kind terms. For instance, a change in usage seems more naturally explained by our coming to mean something different by, say, “chair,” than it is by our discovery that some things we have always called chairs weren’t *really* chairs at all. Meaning is here equated with an additive rather than a dynamic combination of the general characterizations and putative extensions, and any change in either will amount to a change in meaning. This view gets at least some of its plausibility from the problems with the more platonistic picture, but it has serious problems of its own.⁴⁰

Not only does it seem like an unintuitive story for natural kind terms, but it is also by no means obviously true for terms like “chair” either. For instance, a speaker can, being fully aware of his own past usage and the linguistic usage in his community, still argue that some of the items that the community calls “chairs” are, in fact, stools. If he provides good reasons for his claim (the pieces of furniture in question are structurally much closer to prototypical stools than to prototypical chairs, the furniture in question was originally made and marketed by a speaker who was notoriously bad at distinguishing stools from chairs, etc.), then it is not an acceptable reply to his criticism to say, “Listen, they must be chairs, because that’s what everyone in the community calls them.” Indeed, I think that we see something very much like this when certain high-backed bar-stools are assumed to be

³⁹ Davidson, following Bilgrami, claims that “the yearning for such norms is a nostalgic hangover from dependence on a Platonic conception of meaning.” (Davidson 1993a, 145.) See also: “Unless meanings are laid up in heaven, they must depend on the actual practice of someone or some group.” (Davidson 1993b, 117.)

⁴⁰ Still, it should be clear that having all of this information will allow us to ‘communicate’ with the interpretee in the sense of being able to get him to do things, to retrieve information about the world from him, etc.

stools (a natural inference given what they are called), though they are probably more accurately characterized as chairs. Such accounts can't understand content in terms of truth, but rather, just truth as the agent conceives it,⁴¹ and it isn't clear that one can take this latter notion as primitive. After all taking something to be true seems to presuppose the possibility of one's being mistaken, but if what one means is equated with what one takes to be true, such a possibility is ruled out.

However, such pessimism about linguistic norms is unwarranted. Both views discussed above assume presuppose that a practice must be criticized by something outside of it if it is to be criticized at all. However, the 'logical' requirement that one's beliefs and applications be consistent can provide 'internal' grounds for criticizing a practice, and thus will make it possible for one to make non-perceptual errors with one's terms. One can criticize linguistic practice without going 'beyond' it.

One thus need not look beyond our own linguistic practice to insist that what we mean by a term corresponds to how we *should* be (rather than *are*) disposed to apply it. The normativity of the "should" comes from the need to resolve any tensions between *our own* dispositions and beliefs. There is thus nothing mysterious about the norms involved. Lexical norms come merely from consistency constraints upon our attitudes towards a term's general characterization and putative extensions. Lexical norms follow from the norms pertaining to consistency that come from understanding an agent's utterances in terms of truth. Present usage can be criticized by reference to its divergence from an equilibrium that is itself determined by (but not equated with) present usage.

⁴¹ This is fairly explicit in Bilgrami's account, which claims to make no use of notions like truth or reference.

4.2 On sense and meaning

At this point it might be worth discussing how the picture outlined above maps on to the traditional sense/meaning distinction.⁴² (What follows should not, however, be taken as a piece of Frege exegesis). On the account above, we can correlate a term's sense (roughly, its *cognitive* value) with its present putative extension and general characterization, while we can correlate its meaning (its semantic value) with the equilibrium determined by its sense. This allows the sense to be, in a fairly intuitive sense, the mode in which the meaning is presented. It has the truth conditions of a speaker's utterances to be determined by the meanings of his terms, while still allowing the truth of *de dicto* belief ascriptions to be determined by the senses of the terms involved.⁴³ Since equilibria are a function of the terms putative extension and general characterization, there remains a clear sense in which this picture allows that sense determines meaning. Once again, this is not meant to be a piece of Frege exegesis, but is only meant to illustrate how the framework outlined above can capture much of the intuitive distinction that so many have found appealing in Frege's essay.

5. How there might not be any (or any *single*) endpoint

That there might not be anything corresponding to the regulative ideal becomes clear when we consider cases where there seem to be many or no equilibria available. Given a set of putative extensions and general characterizations there will, of course, always be many potential equilibria, since if there is a conflict between the putative extension and the general characterizations, it will always be possible to remove the inconsistency by altering either of the two sides. However, an equilibrium is "accessible" if it does not require giving up elements that are too deeply entrenched, and not every *potential* equilibrium will be an *accessible* one. It should also be noted that how much can be given up will also

⁴² See Frege 1892.

⁴³ In a sense that will be made clear in chapter three, the senses of a speaker's terms will determine what substitutions he will be willing to accept, and thus which *de dicto* attributions will be true of him.

depend, to a certain extent, upon how much must be given up to reach alternative equilibria. If there is a possible equilibrium that requires giving up next to nothing, it can make what would otherwise be accessible equilibria inaccessible. The comparative entrenchment of various elements in the pre-equilibrium state will make various potential equilibria inaccessible. For instance, until quite recently I believed that fish never give birth to live young. It was pointed out to me that guppies (which I took to be fish) do give birth to live young. I could then have made my beliefs consistent in two ways, give up the belief that fish never gave birth to live young or stop calling guppies fish. I would have taken the second of these options to be perverse. Not because there would be something ‘objectively’ wrong about making such a decision, but rather because (like most, if not all, people) I do not have the “don’t give birth to live young” belief nearly as deeply entrenched as the belief that guppies are fish. More often than not, we will feel that one particular way of resolving the inconsistency is the one that gets at what ‘we really meant all along.’ Nevertheless, there are occasionally instances where a number of possible resolutions seem equally accessible (the debate over whether “fish” should have been treated as a natural or functional kind term may have been one of these cases), and there are other cases where there seem to be no acceptable equilibria and either an ‘unperceived ambiguity’ may have to be posited (Field’s discussion of “mass” might be an instance of this), or the term may be given up as unusable.

6. How to understand the dialectical ‘development’

The dialectic outlined above can be understood as describing how we come, over time, to discover what our terms have always meant, as describing how what we mean by our terms comes to change over time, or even as how, over time, we come to mean anything (determinate) at all (with nothing (determinate) actually being meant until the equilibrium is

reached).⁴⁴ For terms not already at an equilibrium (which may be, for all we know, most of our terms) we are left with the question of whether what we (now) mean by them should be identified with the present stage in the dialectic, or where we hope the dialectic to be heading. Which of the two alternative descriptions we should use has been the subject of considerable, and often inconclusive debate.⁴⁵

It is difficult to decide whether such a process should be described as our changing what we mean or as our changing our conception of what we always meant, because “meaning”⁴⁶ is precisely one of those terms for which we have not even come close to reaching any sort of equilibrium. Not only are many of our general intuitions about language and meaning in conflict with each other, but these ideas are often in conflict with the way we actually apply semantic terminology.⁴⁷ The presence of conflicting intuitions and practices associated with our semantic terms gives proponents of the alternative ways of describing linguistic change material to support their claims about whether or not the meaning of our terms has changed through the dialectical development of our linguistic practice.

As a result, a good deal of the contemporary discussion about thought and linguistic content can be usefully understood along the general model of linguistic development outlined above. That is to say, there is, roughly, a set of general beliefs we have relating to meaning and content, and our semantic terms have a putative extension. These do not involve cases of saying “look at that meaning” or “did you see that content over there?” in

⁴⁴ I know of no one who defends the extreme version of this view, though Field’s account of ‘partial reference’ could be understood as a version of the more moderate alternative which understands the dialectic as describing how what we mean becomes more *determinate* over time. (See Field 1973, 1974.)

⁴⁵ This debate has been, of course, most prominent with respect to theoretical terms in science (with writers like Feyerabend (1975) advocating the former type of description and others like Putnam (circa 1975) advocating the latter), but the same issues carry over to other types of terms as well.

⁴⁶ As is the case with “reference,” “belief,” and other such semantically loaded terms.

⁴⁷ See Kay 1983 for the claim that like many other ‘folk’ theories, our tacit theory of language is not globally consistent and incorporates conflicting conceptions of what language is.

the way that the putative extension of “cat” would depend upon our utterances of “that’s a cat,” “look at that cat” or “did you see that cat over there?” Rather they involve our ‘standard’ or ‘commonsense’ ascriptions of content, and similar evaluations of the meanings or truth-conditions of the statements of others.⁴⁸ Much of the most interesting work in the philosophy of mind and language has arisen from the gradual realization of just how much tension there is between our general beliefs about meanings and contents and our standard practices of ascribing them.

If the general picture of linguistic development outlined above is right, and it also applies to our theorizing about ascriptions of meaning and content, then one should not expect knock-down arguments that will leave everyone, if not happy, at least convinced. Thinking that the debate will necessarily have this character is the result of thinking that there is some antecedent truth about meaning and content and that we just have to sniff around enough until we find it (i.e., some form of the platonist picture.) Resolving the tensions between our ascriptional practices and theoretical commitments may involve some give and take, and it thus may involve giving up things that would be, all things being equal, worth having. (Again, the parallel with ethical/practical deliberation is worth noting.) Because of this, it is important not to motivate a particular view by something like the argument from elimination (i.e.: by pointing out that all the other positions force us to give up something which we should reasonably want), since it seems quite probable that there is no position that would not force us to give up something which we would reasonably want.

Furthermore, when evaluating these conflicts between the ascriptional practices and general beliefs, *all* the general beliefs should be taken into consideration, especially since the tension between the ascriptional practices and the general belief may be the result of

⁴⁸ For a similar appeal to content ascriptions to fill such a role, see Devitt 1994, 1996. Of course, this notion of just what the ‘natural’ ascriptions of content are (or even the notion that there is such a thing) is not without its problems.

tensions among the general beliefs themselves.⁴⁹ Indeed, this process of tension resolution allows us to look at things on, as it were, a more global level. That is to say, when resolving a tension in one area, general beliefs we have in other areas can be taken into consideration. This looking towards a ‘total package’ of views can considerably alter how things work out at a local level.

In the philosophy of language, which often considers itself the ‘foundational’ part of philosophy,⁵⁰ one can occasionally get the impression that it is ‘not fair’ to bring up considerations from other areas of philosophy when evaluating a semantic theory. For instance, while a theory of meaning might be rejected if it entailed that we could never learn or know a language, if it entailed that we could never know anything about, say, the past, or even the world around us, such a result might instead just be considered an interesting (if disturbing) consequence.⁵¹ This attitude isn’t really defensible if the general view presented above about what is involved in settling issues about meaning is right. If one’s account of content entails that we can’t know anything about things we intuitively take ourselves to know, then this counts as an at least *prima facie* reason to think that one’s account is wrong.

7. Some general beliefs about meaning

There are a number of general beliefs we have about meaning that non-philosophers rarely attempt to formulate explicitly, and not much should be taken to hang upon the

⁴⁹ The fact that the general beliefs that “fish” picked out a natural and a functional kind were inconsistent was made manifest by the discovery that whales, while of the same ‘functional’ kind as other fish, were not of the same ‘natural’ one.

⁵⁰ “The theory of meaning is the fundamental part of philosophy which underlies all others.” (Dummett 1973, 669.)

⁵¹ The fact that the classical theory of ideas was held so long in spite of its clearly disastrous consequences for epistemology might be an instance of this. Michael Dummett’s work could also be understood this way, and the sentiment is aptly expressed, though not endorsed, by Fodor and LePore in the following:

Serious though these worries are, however, they are all, as it were, *external* to functional role semantics as such; they impugn its consequences for epistemology and ontology rather than its coherence. And a convinced New Testament Semanticist might be prepared to bite the bullet. If the doctrine that meaning supervenes on intra-linguistic relations has relativistic, idealistic or even solipsistic consequences, then perhaps we had better learn to live with relativism, idealism or even solipsism. (Fodor and LePore 1993, 20.)

particular formulation of them here. For instance, it seems clear that we take our words to refer to (and our thoughts to be about), independent objects which they represent to a greater or lesser degree of accuracy. This ‘belief’⁵² is fairly deeply entrenched, and will be formulated for present purposes as:

(1) Our beliefs and sayings are true or false of (about) the world around us.

There are a number of ‘philosophical’ beliefs about meaning and content which while perhaps not widely shared, owe their plausibility to their status as either explications of (1), or things that must be true if (1) is to be.⁵³

Another deeply entrenched belief has to do with what is often referred to as “self-knowledge,” “first-person authority,” “privileged access” etc. We tend to know better than others do, and seem to know in a different (“unmediated”) way, what we believe. That is to say:

(2) We tend to know what we are thinking.⁵⁴

A further belief that is equally entrenched is that our behavior can be explained by our beliefs and desires. Once again, while this belief is often not explicitly formulated, it is clear that we take an agent’s beliefs to be relevant to his behavior. For the time being the belief can be expressed as:

(3) Our behavior is explained by our beliefs and desires.⁵⁵

⁵² While formulated here as explicit beliefs, some of these ‘beliefs’ might better be thought of as part of a, possibly pre-intentional, “Background” of abilities and disposition that can have a lot to do with what type of equilibrium can be reached for these terms. (For a fuller discussion of this, see appendix F.)

⁵³ For instance, the belief that we refer to objects in virtue of being able to specify (through demonstratives, descriptions etc.) what those objects are can be understood as an (unsuccessful) attempt to explain how we refer to items in the world. The belief that our terms have determinate extensions is related to (1) and may even be a relatively philosophical explication of it.⁵⁴ Like (1), there is no shortage of philosophical beliefs trying to make sense of (2).⁵⁵ This should not, of course, be taken as an empirical generalization about what we actually say when we explain someone’s behavior. There is, of course, considerable debate about how (3) should be explicated, in particular, about what type of explanation is involved (causal, lawlike, hermeneutic, etc.).

Other beliefs that could be seen as part of the general characterizations of our semantic terms are the following:

- (4) When an agent sincerely utters a sentence with a certain content, he has a belief with the content of his utterance.⁵⁶
- (5) There is a fact of the matter about what we mean.⁵⁷
- (6) We speak a shared language.⁵⁸

While it is undoubtedly the case that these more philosophical characterizations of semantic content are more easily given up when they conflict with our ascriptional practices, it should not be assumed that this is always so. Since the philosophical characterizations are often the product of reflection upon (and taken to be an explication of) the more primitive beliefs, giving up some of the philosophical beliefs may require giving up some of the primitive ones (or at least coming up with a new understanding of them).

For instance, methodological solipsism⁵⁹ (roughly the idea that the meanings of our words and the contents of our thoughts could be characterized independently of our relations to items and events outside of our heads) may seem like a comparatively recent philosophical prejudice about meaning that we have no reason to keep it if it conflicts with our ascriptional practices. It provides, however, a fairly elegant way of accounting for many of the more primitive beliefs.⁶⁰ It seems initially to do well with behavioral explanation, the factuality of meaning, self-knowledge, and to provide an account both of how language relates to the world, and of how we are related to language. When methodological solipsism was given up, many of these more primitive beliefs were questioned (we don't really know what we mean, etc.), and, indeed, it became quite popular to suggest that there is a methodologically solipsistic 'narrow' content that is, in

⁵⁶ This belief and various attempts to explicate it are the subject of chapter two.

⁵⁷ Given how one feels about facts, this can lead to the more obviously philosophical belief, "Semantics is an empirical science like any other" (Devitt 1994, 545).

⁵⁸ Which has associated with it more philosophical beliefs such as that meanings are necessarily public, and that meaning depends in part upon convention.

⁵⁹ A position defended most vigorously in Fodor 1981, 1987. ⁶⁰ Though not all. In particular, the idea that we speak a shared language drops out.

some sense, independent of our ascriptional practices and used to make true the general beliefs.⁶¹

8. Attitude ascriptions

It was suggested above that attitude ascriptions can be treated as evidence for the conditions of application of our semantic terms. The role of ‘thought experiments’ of the sort made famous by Kripke, Putnam, and Burge can be understood as a way of eliciting such data about how semantic terms are actually applied. Understood this way, their thought experiments are comparatively unambitious: they merely draw our attention to the type of ascriptions we are inclined to make, but they do not commit us to the truth of such ascriptions (they relate to what we *do* say in such situations rather than what we *should* say). Of course the authors believe that what we do say is, in fact, what we should say in these cases, but even those who do not share their intuitions about the thought experiments can agree that most people will attribute content the way the authors suggest.

It should also be stressed that, when saying that the ascriptions themselves are fairly uncontroversial, one is saying that people faced with either of the two possibilities that, say, Burge describes will ascribe content in the way Burge says they would. It does not involve saying that most people, when faced with a thought experiment involving *both* possibilities, will describe the pair the way Burge does. Being faced with such contrasting pairs may cause us to bring our general beliefs into play, and in such cases, our ascriptions are more likely to reflect our general beliefs about content than our general ascriptional practices. That is to say, it does seem true that people when faced with the

⁶¹ Some more philosophical beliefs relating to our semantic terms that will come up in the course of the dissertation are:

- i. Russell’s principle: “a subject cannot make a judgment about something unless he knows which object his judgment is about.” (Evans 1982, 89.)
- ii: The Unity Constraint: The contents that are evaluated as true or false are the very same contents that explain behavior. (Bilgrami 1992.)
- iii: If the speaker intends to use a name 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. (Kripke 1972.)

arthritis/tharthritis cases separately will ascribe different contents to Bert and his counterpart, but it need not follow that people who are faced with both Bert and his counterpart will say that they mean anything different by “arthritis.”

Four writers who have pointed out how certain aspects of our ascriptional practice were in tension with general beliefs we had about meaning are Saul Kripke, Hilary Putnam, Tyler Burge and Mark Wilson. They do not necessarily present themselves as doing this, but understanding them in this way does, I believe, no violence to their work. Still, this is a revisionary history of what is going on, not an exposition of what the four authors took themselves to be doing. While these writers all present us with examples of how our ascriptional practices can conflict with general beliefs about meaning, the ascriptional practices become increasingly easy to give up as one moves from Kripke to Wilson (or, correspondingly, the general beliefs become progressively harder to give up). Indeed, it is not surprising that Burge’s and Putnam’s thought experiments are the ones to have generated the most controversy. In the Kripke cases, we stick too closely to our ascriptional practices, in the Wilson cases, too closely to our general beliefs for controversy easily to get off the ground.

As stressed in section six, in evaluating how these tensions are to be resolved, it is important to get a complete picture of what is involved. Resolving tensions in a certain way locally may turn out to be incompatible with the most attractive global resolution. This involves first getting a thorough account of what the ascriptional practices are, since they often get much of their force collectively. For instance, while the ascriptional practices described by Burge may initially seem to conflict too much with our general beliefs to be accepted, once these general beliefs are modified to account for the Kripke and Putnam cases, the degree of conflict may be considerably reduced. 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. Indeed, I would guess that before 1972, Burge’s thought experiment would have been something of

a non-starter. It would have been too natural for philosophers to say Bert just meant by “arthritis” “a disease of the joints and thigh.” It was only because such descriptive theories had been discredited already that such a response came to be a much less appealing option. In much the same way, sticking with the ascriptional practices in the Wilson cases would not seem even a remote possibility if one were not familiar with the Burge cases, which are, in many respects, structurally similar.

Each of the four authors highlights aspects of our ascriptional practice that suggest that some set of criteria that might have been assumed adequate for determining reference or content, can, in fact, underdetermine and even misidentify what we intuitively take ourselves to be referring to. In order to accommodate the relevant attributions, other factors -- such as the speakers’ linguistic history, physical environment, social environment and even future history of use -- must be incorporated into the proposed set of criteria. The following will highlight the types of underdetermination and misidentification these authors bring to our attention.

8.1 Kripke: proper names and history of use

The first, and in some sense paradigm case of such a conflict is Kripke’s discussion of proper names. The bit of our ascriptional practice he pointed to was our habit of interpreting people who use a proper name as referring to the historical source of most of the information associated with that name.⁶² The general belief that Kripke explicitly targeted was the idea that the reference of a speaker’s use of a name was determined by the descriptions he associated with that name, though little would have changed if other sorts of recognitional capacities were added as well. The types of problems Kripke raises against the descriptive theory of names (that the ‘internal’ information can often underdetermine or even misidentify the referent) will turn out to apply to internalist accounts of many other terms as well.

⁶² This particular way of formulating just what Kripke pointed to owes more to Evans’ discussion of Kripke (in Evans 1973) than to Kripke’s own discussion.

Underdetermination: Feynman

The first of these problems relates to the fact that, while it may be the case that, for some names, we can come up with a set of descriptions which, in fact, is satisfied uniquely by the intended referent, the descriptions we associate with a name often radically underdetermine who that name refers to. Kripke gives an extended example of this sort of underdetermination in his discussion of our term “Cicero,” but the general point is summed up quite nicely in the following discussion of “Feynman.”

Consider Richard Feynman, to whom many of us are able to refer. He is a leading contemporary theoretical physicist. Everyone *here* (I’m sure!) can state the contents of one of Feynman’s theories so as to differentiate him from Gell-Mann. However, the man in the street, not possessing these abilities, may still use the name ‘Feynman.’ When asked he will say: well he’s a physicist or something. He may not think that this picks out anyone uniquely. I still think he uses the name ‘Feynman’ as a name for Feynman.⁶³

The basic point here should be clear. Kripke draws our attention to the fact that we are typically taken to be speaking about people we have no way, if left to our own resources (i.e., we cannot help ourselves to expertise of other speakers, written information, etc.), of uniquely identifying (other than using the name itself).

Misidentification: Einstein, Thales

The second problem is even more serious than the first. Not only is it the case that the descriptions a speaker associates with a name may fail to pick out its referent uniquely, these descriptions may actually pick out someone other than the person who the term is intuitively taken to refer to. This can be seen in Kripke’s more extended discussion of names like “Gödel,” “Aristotle,” but is summed up quite clearly in the following paragraph:

Even worse misconceptions, of course, occur to the layman. In a previous example I supposed people to identify Einstein by reference to his work on relativity. Actually, I often used to hear that Einstein’s most famous achievement was the invention of the atomic bomb. So when we refer to Einstein, we refer to the inventor of the atomic bomb. But this is not so. Columbus was the first man to realize that the earth was round. He was also the first European to land in the western hemisphere. Probably none of these things are true, and therefore, when people use the term ‘Columbus’ they really refer to

⁶³ Kripke 1972, 81.

some Greek if they use the roundness of the earth, or to some Norseman, perhaps, if they use the 'discovery of America.' But they don't.⁶⁴

Such cases convinced most philosophers that the descriptive theory of names was deeply flawed, and it seemed that other internalist accounts of proper name meaning would do no better.⁶⁵ What it was about us that seemed to relate us to the bearers of the names we used was not just the descriptions or the discriminating capacities we had, but also our being part of historical chains involving the use of each particular word.⁶⁶ Of course it is always open to the defender of the descriptive theory of names to be hard-headed (the no counterexamples, only consequences approach), and claim that we should just understand this aspect of our ascriptional practice as confused, but it is notable that, to my knowledge, no one seems to take this line.

8.2 Putnam: natural kinds and the physical environment

In much the same way, Putnam's Twin-Earth stories can be understood as showing that our ascriptional practices involving natural kind terms are in conflict with the idea that what we mean supervenes on anything like our brain states.

Underdetermination: beech, elm

Putnam's paper is full of examples of underdetermination, and one of the upshots of his Twin-Earth examples is that underdetermination is, at least in principle, possible even for familiar terms such as "water" or "gold," though he gives a more 'down to earth' example with his own use of "beech" and "elm":

Suppose you are like me and cannot tell an elm from a beech tree. We still say that the extension of 'elm' in my idiolect is the same as the extension of 'elm' in anyone else's, viz., the set of all elm trees, and that the set of all beech trees is the extension of 'beech' in *both* of our idiolects. Thus 'elm' in my idiolect has a different extension from 'beech' in your idiolect (as it should).⁶⁷

⁶⁴ Kripke 1972, 85. Much the same point is made by Donnellan, who, even if it were the case that the only description we associate with "Thales" is "the Greek philosopher who held that all is water." (Donnellan 1972, 374.)⁶⁵ Though, for a dissenting opinion, see Searle 1983.

⁶⁶ For another expression of this idea, see Geach 1972, 155.

⁶⁷ Putnam 1975, 226.

Once again, it seems that we treat people as referring to things in the world independently of their having any descriptions or discriminative capacity that would allow them to pick out the objects in question.

Misidentification: unicorns

Cases where the descriptions associated with a natural kind actually fail to be true of the kind in question but are true of a different kind do not appear in Putnam's paper, but, given what has been laid out so far, they are not hard to construct. Indeed, Kripke's discussion of "unicorn" provides such a case:

The mere discovery of animals with the properties attributed to unicorns in the myth would by no means show that these were the animals that the myths were about: perhaps the myths were spun out of whole cloth, and the fact that animals with the same appearance actually existed was mere coincidence. In that case, we cannot say that the unicorns of myth really existed; we must also establish a historical connection that shows that the myth was *about* those animals.⁶⁸

Further examples are easy enough to construct, and the general point remains clear. We tend to treat the physical kinds that we are historically connected with as affecting what we mean by our words in ways that cannot be accounted for within an descriptivist framework.

Putnam's examples are another instance where the general belief has tended to cave in to the ascriptional practice, but it is noticeable that in this case there has been considerably more resistance. Not only have there been many attempts to have it both ways (narrow content matches general beliefs, wide content matches ascriptional practice, etc.),⁶⁹ but there have also been those who argue that the ascriptional practice of saying that Caesar and his Twin-Earth counterpart meant something different by "water" is simply mistaken, and should be revised in light of our general beliefs.⁷⁰

⁶⁸ Kripke 1972, 157.

⁶⁹ This reaction was, of course, made to Kripke's work too.

⁷⁰ It should be noted that the Putnam and Kripke examples paved the way for a fairly enthusiastic reworking of the general beliefs about meaning in terms of things like "causation." One can thus see the emergence of a new framework that may conflict with other aspects of our ascriptional practice. For instance, our 'natural' ascriptional practices would be to ascribe contentful thoughts and sayings to Davidson's 'Swampman,' and while most of the individualist and internalist frameworks can make sense of this practice, the causal/historical framework cannot. As a result, Davidson argues that people who attribute thoughts to the Swampman are confused (see Davidson 1987), while others have taken

8.3 Burge: non-natural kinds and the social environment

Tyler Burge's much discussed example of Bert and his use of "arthritis" suggests that our ascriptional practice is in conflict with the idea that the contents of our thoughts and the meanings of our words are, in principle, independent of the contents of the thoughts and meanings of words of those around us.⁷¹

Misidentification: arthritis

While the sections on proper names and natural kinds began with underdetermination and moved on to misidentification, this one will do the opposite. This is partly because Burge's thought experiments usually involve cases of misidentification, and this should not be surprising. If one or two of the beliefs I associate with a given name turn out to be false, it is quite unlikely that there will be a person out there who does satisfy all the descriptions associated with the name. On the other hand, if some of my beliefs about, say, chairs happen to be false (for instance, I believe that they must have four legs and be made of wood), there will still be a recognizable class of objects about which all my 'chair-beliefs' will be true.

By far the most famous case of such misidentification involves Bert and his claim that he has "arthritis" in his thigh.⁷² Bert has a number of correct beliefs about arthritis such as that he has had arthritis for years, that older people are more likely to have it, that stiffening joints is a symptom of arthritis, etc., but he also has the false belief that his arthritis has spread to his thigh. When he tells this to his doctor, the doctor tells him that arthritis is specifically an inflammation of joints, so Bert gives up his belief and goes on to wonder

the Swampman to show that there is something deeply wrong with the general beliefs about meaning that make up the causal/historical framework (see Fodor 1994).

⁷¹ That is to say, while the contents of our thoughts etc. are *causally* dependent upon the contents of those around us (since it is our interaction with them that causes us to have the properties that constitute our contents), there is no *constitutive* relation between our contents and theirs.

⁷² See Burge 1979, 77 and following.

what could be wrong with his thigh. In spite of his false belief, Bert is still taken to refer to *arthritis* by “arthritis.”

Burge now asks us to consider a counterfactual situation in which Bert “proceeds from birth through the same course of physical events that he actually does, right to and including the time at which he first reports his fear to his doctor,”⁷³ the only difference being that, in the counterfactual situation, Bert’s society uses the word “arthritis” to pick out inflammations outside of the joints as well, particularly in the thigh (this cluster of inflammations will be referred to as “*tharthritis*”). In such a case, Bert would be taken as expressing a true belief when he said “I have arthritis in my thigh,” and that is because he taken to be referring to *tharthritis* by his word “arthritis.”

These cases suggest that what we take Bert to mean by the word “arthritis” is determined not just by facts about his internal make up (or such facts plus facts about the purely causal relations between him and his environment), but also by the use of the term in the society to which he belongs.

Underdetermination: arthritis again

Burge does not really focus on cases that involve underdetermination, but given the scenario already in play, such a case is easy enough to construct. What distinguishes *tharthritis* from arthritis is that one can get the former, but not the latter, in one’s thigh, so a case of underdetermination can be seen if we consider a speaker much like Bert but with the qualification that he has absolutely no opinion about whether or not one can get “arthritis” in one’s thigh. His beliefs would thus match equally well with either arthritis or *tharthritis*, but in each case one would attribute to him the concept used by the surrounding community (indeed, the impulse to do this is much stronger in the cases of underdetermination than it is in cases of misidentification).

⁷³ Burge 1979, 77.

Much more so than with the Putnam and especially the Kripke cases, people⁷⁴ when faced with the practices Burge appeals to are willing to simply stick by the general belief and say that this part of our ascriptional practice is ‘confused’ or ‘sloppy.’

8.4 Wilson: future usage

Finally, there is a range of cases brought to our attention by Mark Wilson that suggest that our everyday ascriptions of content not only reflect a sensitivity to the agent’s internal make up, and the structure of his physical and social environment (including the past history of use of the terms he is using), but also reflect ‘accidental’ developments in the term’s use in the time subsequent to any particular utterance.

Underdetermination: Grant’s Zebra

For instance, around 1820 a fellow named Grant ‘discovered’ a purported species of Zebra native to Kenya and called it “Grant’s Zebra.” A few years later, the term “Chapman’s Zebra” was introduced for a morphologically distinct type of Zebra found in what was then Rhodesia. Later still it was discovered that the two purported species of Zebra interbred near the Zambezi river and that, morphologically, one gradually faded into the other. Grant was then taken to have discovered a subspecies of the larger species *Equus burchilli*. This story is typical of the way we describe such cases. However, this aspect of our ascriptional practice is actually at odds with the general belief we have that what we mean by our words is independent of accidental features of our society’s *future* linguistic practice. This general belief is in conflict with our practice in the Grant’s Zebra case because, on reflection, it seems to be merely a *historical accident* that the term “Grant’s Zebra” came to have the extension it did. If the taxonomists had investigated the area around the Zambezi river *before* they hit deepest Rhodesia, then they probably would have ‘discovered’ that Grant’s Zebra could be found through most of East-Africa gradually changing into a different subspecies as it drifted towards Zimbabwe. In such a case, “Grant’s Zebra” would have picked out the entire species, not just the subspecies found in

⁷⁴ For instance, Davidson, Bilgrami, Camp, my long past self.

Kenya. This is, of course, just a single example, but Wilson maintains that similar examples can be found “in virtually every case of enlargement of our world view through scientific progress.”⁷⁵

Misidentification: natural and functional kinds

It seems clear that many of Putnam’s remarks about natural kinds could be understood as saying that a society’s present usage can misidentify what a term refers to because there can be a diachronic as well as a synchronic division of linguistic labor. However, the fact that the division of labor can stretch out diachronically need not violate the belief that what a term means is determined by contemporaneous linguistic practice. Our present practice may make it clear that a term is a natural kind term and what constitutes being a natural kind, and if this is the case, how we would react to, say, certain chemical discoveries in the future may already be determined. The cases Wilson has in mind are more radical. In these cases it is the subsequent history itself that determines whether a term is a natural rather than a functional kind term, and so can also determine that large ranges of previous application (which would have been correct if the term had turned out to be a functional kind term) were incorrect. For instance, when the term “gold” (or its equivalent) was introduced, the introducers might have thought (to the extent that they thought about it at all) that the term picked out both a natural and a functional kind. With the discovery that the class of ‘gold-looking’ things contained more than one kind, one had the choice of giving up either the belief that ‘gold’ picked out a natural kind, or the belief that it picked out a functional one. The tradition of which we are a part chose to preserve the belief that it was a natural kind, and so early members of this tradition who called iron pyrite “gold” can be taken to be mistaken, even if they had no way of distinguishing the two. Had the tradition gone on to treat “gold” as a functional kind, the very same speaker would have

⁷⁵ Wilson 1982, 572. ⁷⁶ This would, I believe, be the upshot of Field’s claims about the ‘Newtonian’ use of the term (see, for instance, Field 1973).

said something true (though he might have had the false belief that “gold” picked out a naturally uniform kind).

Nevertheless, the conflict we see in the Wilson cases is one in which many, if not most, people (possibly including Wilson himself) are willing to say that the ascriptional practice is mistaken and hold on to the general belief.

9. The possibility of many equilibrium points: indeterminacy

One hopes, of course, that a systematic investigation of our ascriptional practices and the beliefs we associate with our semantic concepts will converge towards a stable interpretation of them. However, there is no guarantee that it will be possible to find any such convergence, or, even if it is, that there will be a *single* definitive way for the various elements to converge. Just as there may not be a single accessible equilibrium between the putative extensions and general characterizations of non-semantic terms such as, say, our old use of “fish” or “mass” (the former seeming to have at least two accessible equilibria, the latter quite possibly having no accessible equilibria),⁷⁶ there may be no single accessible equilibria between the putative extension and general characterizations for any of our semantic terms either. In particular, the answer to the question of whether, when certain bits of our practice conflict with our general beliefs, we should revise our beliefs or the practice itself depends upon our attitude towards the centrality of the attributions or beliefs in question. For instance, one’s attitude as to whether or not whales are fish depends upon how central one takes the assumption that fish are a natural kind to be. If one takes it to be fairly central (as scientists are inclined to do), then it will seem ‘obvious’ that the best equilibrium will involve giving up the practice of calling whales fish. On the other hand, if one takes the assumption about natural kinds to be comparatively unimportant (as fishermen are inclined to do), then it may seem equally clear that one should reach an equilibrium in which whales *are* fish. In much the same way, our attitudes towards the comparative centrality of various semantic intuitions will be influenced by whether we are

interested in, say, cognitive science (naturalistic intuitions about meaning presupposed by the program will, of course, be fairly central) or epistemology (reductive/internalist intuitions may seem to be of little concern if they commit one to skepticism etc.). When such disputes occur, there may be no more reason to think that one of the competing semantic theorists is ‘deeply confused about the nature of meaning’ than there is to think that one of either the scientist or the fisherman must be ‘deeply confused about the nature of fish.’ Still, people who accept one equilibrium will have a tendency to view those who adopt another equilibrium as confused.⁷⁷ When such accusations are not made and each party accepts the usage of the other as a genuine alternative, there is still a tendency to treat one’s own use as the norm. Scientists will contrast what fish are, ‘strictly speaking,’ with “fish” in its ‘vulgar’ sense, while laymen will contrast the ‘normal’ use of “fish” with its ‘technical’ use.

9.1 ‘Quinean’ and ‘Kuhnian’ indeterminacy

Theorists who assign different weights to various general presuppositions about meaning and content may simply have ‘incommensurable’ views about how a coherent position should be reached, and this may produce genuine indeterminacy about how the contents of our thoughts and utterances should be properly individuated. That is to say, different interpreters, both fully aware of all the physical facts, might still legitimately disagree about how to interpret a given speaker. This possibility brings to mind, of course, Quine’s thesis of the indeterminacy of translation, which he characterized as follows:

Manuals for translating one language into another can be set up in divergent ways, all compatible with the totality of speech dispositions, yet incompatible with each other. In countless places they will diverge in giving, as their respective translations of a sentence of the one language, sentences of the other language which stand to each other in no plausible sense of equivalence however loose.⁷⁸

Quine found translation to be indeterminate because all the relevant evidence underdetermined what the correct interpretations of a speaker’s words were. The position

⁷⁷ As Ishmael’s former shipmates put it, the scientist’s reasons were “altogether insufficient,” and, indeed, “humbug.”

⁷⁸ Quine 1960, 27.

outlined above suggests that translation may (also) be indeterminate because what counts as the relevant evidence is *itself* indeterminate. Different interpreters, each fully informed about the facts about word usage in past, present, future and even counterfactual situations, could still disagree as to how these facts should be interpreted because they had different preconceptions about meaning and thus about what sort of evidence was relevant to an assignment of meaning, what the criteria for change of meaning were, etc.⁷⁹

For instance, if two interpreters had decided to come up with manuals for my early use of “vegetable,” one could quite easily imagine the following conflict emerging: the first interprets me as having the false belief that mushrooms are vegetables, while the second interprets me as meaning something different by “vegetable” than most of my peers. The first will interpret me as having meant the same thing by “vegetable” during my entire life, while the second will suppose that the meaning of this term changed for me when I stopped applying it to mushrooms. The first may give primary importance to being faithful to my own conception of what was happening (I took myself to have had a false belief about

⁷⁹ Some (see, Hylton 1982, 180) have argued that something of this sort was what Quine had in mind with his indeterminacy thesis, but Robert Kirk has argued that treating indeterminacy this way trivializes it:

Exactness of literal equivalence by no means entails exactness of cultural equivalence. But Quinean indeterminacy would be a trivial thing if it depended on the fact that the constructors of different translation manuals may have divergent purposes and priorities. Not only would it be trivial, it would be quite useless for Quine because it would be consistent with the most extreme Platonism about meaning. For Platonists can easily agree that when translators have different purposes they will want to associate sentences with different Platonic meanings. If we are not to misrepresent the indeterminacy thesis, therefore, we must assume our translators share the same purposes and priorities, whatever these may be.... it seems best to assume that these purposes and priorities are scientific. The translators are to be supposed to be aiming at exact literal synonymy -- whatever that means exactly. (Kirk 1986, 73.)

Kirk’s criticism seems, however, to miss the mark, and his claim that indeterminacy stemming from interpreters having different purposes and priorities “is entirely consistent with extreme Platonist theories of meaning, and has no philosophical interest” seems to be based upon trivializing what these different purposes and priorities could be. Admittedly, if the differences all resulted because one interpreter was interested in preserving the “cultural” meaning of what is said while the other was interested in the “literal” meaning, the thesis would be of comparatively little interest. However, in the cases described above, all parties concerned are looking for the ‘literal meaning’ -- it’s just that their “different criteria for the adequacy of a translation manual” cause them to disagree about what is relevant for determining the *literal* meaning. Kirk claims that “the translators are supposed to be aiming at exact literal synonymy -- whatever that means exactly,” but whether there is anything which we exactly mean by “literal synonymy” is precisely what is at issue, and the advocate of Kuhnian indeterminacy may argue that there is nothing which we mean exactly by such a term.

mushrooms), while the second puts greater weight on, say, closely tying what I mean by my terms to those things I causally interact with. The first favors interpretations that attribute shared meanings to members of the same linguistic community, the second favors interpretations that maximize the amount of true beliefs on the interpretee's part. The two theorists may each see the value of the other's desiderata, but disagree over the comparative weight that should be placed upon them.⁸⁰

That there could be such conflicts between interpreters should be clear to anyone with even a passing acquaintance with much recent work in the philosophy of language. Not only is the example structurally similar to that of Bert's use of "arthritis" (with Burge playing the role of the first interpreter, and, say, Davidson playing the role of the second),⁸¹ but other examples of how philosophers differ over what criteria an interpretation should satisfy are also not hard to find. To mention just one such case, some philosophers would more or less discount any interpretation that did not seem as if it could be given some sort of naturalistic grounding,⁸² while others would find this proposed criterion not only dispensable, but also in conflict with fairly central beliefs about the normative nature of semantic concepts. That this sort of conflict could produce 'incommensurable' commitments to divergent manuals has considerable plausibility.

There is little reason not to understand such disputes as instances of indeterminacy.⁸³ What Quine wanted, after all, was a case where "two translators might develop independent manuals of translation, both of them compatible with all speech behavior, and yet one manual would offer translations that the other translator would reject."⁸⁴ Davidson's and

⁸⁰ It should also be noted that this type of indeterminacy does not seem to conflict with the apparent determinacy we have from the 'first person point of view.' While I may be certain that I *never* meant "undetached rabbit part" by "rabbit," I have no such firmness of faith about what I meant by "vegetable" (indeed, I've changed my mind about it on more than one occasion).

⁸¹ See, for instance, Burge 1979, Davidson 1987, 449, and Davidson 1994a, 6.

⁸² For a recent expression of this view, see Devitt 1994.

⁸³ Though, of course, there is no guarantee that they will be. It is certainly *possible* that, say, Davidson and Burge have both worked out stable equilibrium points for our semantic terms, but whether this is, *in fact*, the case is something that will have to be investigated in what follows.

⁸⁴ Quine 1979, 167.

Burge's manuals clearly seem to fit the bill,⁸⁵ and there is no reason to think that any further utterances on Bert's part would settle the issue between them.⁸⁶ When the layman's use disagrees with that of the expert, Burge treats it as a genuine disagreement, while Davidson treats it as a verbal one.⁸⁷

However, while the considerations raised above may suggest that examples of (and a motivation for) indeterminacy may not be difficult to find, it should be clear that the examples and motivation involved will be significantly different from Quine's. Indeed, one might draw a distinction between 'Quinean' and 'Kuhnian' indeterminacy. The former assumes agreement upon what counts as the relevant evidence and how it should be weighted and argues that such agreed upon evidence supports a number of conflicting theories equally well.⁸⁸ The latter is structurally similar to how our choice of scientific theories can be underdetermined due to 'incommensurable' differences between theorists about what counts as relevant evidence etc., and about how such evidence should be weighted.

The two types of indeterminacy need not be mutually exclusive. However, while there is little in Quine that positively rules out the possibility of Kuhnian indeterminacy, his positive views on the constraints on translation and his apparent belief that we couldn't find any actual instances of indeterminacy suggest that he does not have anything like Kuhnian indeterminacy in mind.

⁸⁵ Furthermore, one might want to claim that in the case of disputes about meaning of the sort Burge and Davidson would find themselves in, "either manual could be useful, but as to which was right and which was wrong there was no fact of the matter." (Quine 1979, 167.)

⁸⁶ This might strike some as a controversial claim, but unless Bert is unusually reflective about semantic issues pertaining to his own language, this will probably be true. Of course Bert could always *come* to be reflective in this way, but it would be open to either disputant to claim that such reflection had caused Bert to change what he meant and thus had no bearing on the interpretation of the original utterance.

⁸⁷ This sort of disagreement fits well, of course, with the point in "Two Dogmas" about there not being a firm distinction between differences in belief and differences in meaning once we give up the analytic/synthetic distinction. (See Quine 1953.) It is also in keeping with Davidson's claim that "when others think differently from us, no general principle ... can force us to decide that our difference lies in our beliefs rather than our concepts." (Davidson 1974b, 197.)

⁸⁸ See Kuhn 1973, 326. ⁸⁹ Indeed, Quine seems to have resigned himself to our being unable to discover any examples of the indeterminacy of translation proper, claiming instead that "full holophrastic indeterminacy of translation draws too broadly on language to admit of factual illustration." (Quine 1990, 50-51.)

The nature of Kuhnian indeterminacy also makes it clear why Quine himself had so much trouble coming up with an instance of the indeterminacy of translation proper.⁸⁹ Conflicting translation manuals, like conflicting accounts of the world's physical structure, are likely to be created by people working with conflicting theoretical or methodological commitments.⁹⁰ These commitments cause them to view the phenomena in different ways, as more or less important, etc. Quine, however, for all his talk about the rival manuals, never seemed to doubt that there was a single true 'theory of meaning' (namely, some version of the type of theory one found in *Word and Object*) which settled all the questions about what sort of evidence should be taken into account and what significance it should be given.⁹¹ As a result, Quine's own methodology implicitly denies that the Kuhnian type of indeterminacy holds, and thus makes it extremely difficult for him to find an example for his own thesis.

9.2 Conventionalism

As discussed previously, when scientists discovered that they could not count whales as fish fishermen were not impressed. Nevertheless, while the two parties may have had some inclination to treat the others as 'confused about fish,' it seems quite clear that both ways of categorizing fish were recognizable extensions of the original "fish" practice, and the term could quite easily come to have two meanings (an 'everyday' and a 'technical' sense). If this option is available, then any dispute between the scientists and the fishermen would be a purely verbal one. Each pick out an equilibrium accessible from the previous

⁹⁰ This might also help explain why, after two decades of appearing to swallow the, on the face of it, incredibly implausible claim that there was no fact of the matter of whether we meant *rabbit* or *rabbit-part* by "rabbit," philosophers of language such as Davidson have so easily found themselves treating the question of whether content was individuated individualistically or not as a factual one. Gavagai-driven disagreements can arise between people who share the same general methodology, and are thus unable to discount their rival's manual. When the two interpreters do not agree on methodology, it is much easier for each interpreter to see the other's manual as being the product of 'conceptual confusion,' 'insensitivity to the real nature of psychological explanation,' 'lack of theoretical sophistication,' etc.

⁹¹ "In psychology one may or may not be a behaviorist, but in linguistics one has no choice." (Quine 1987, 5.)

practice. It is simply a matter of convention which of the two we choose. Indeed, we might be able to freely switch from one to the other as we choose (as we seem to do with “weight,” sometimes using it in its technical sense, and other times using it as more or less equivalent to “mass”).

In much the same way, the viability of treating disputes between, say, Davidson and Burge as instances of indeterminacy may lead to a position where it is just a matter of convention how we should individuate content, with the consequence that different groups can do so in whichever way best suits their purposes. Just as the fisherman can work with a ‘functional’ conception of fish, while the scientist works with a ‘biological’ one, the computational linguist could work with an internalist conception of content, while the epistemologist works with an externalist one. Or, once again, the same person might use one while we are doing epistemology, another while doing AI, etc.⁹²

However, the fact that our pre-theoretical intuitions about meaning pull us in many different directions by no means guarantees that there will be equilibria corresponding to each of these directions, and the fact that it may initially seem that there will be a stable equilibrium attached to some understanding of meaning by no means guarantees that there will be such an equilibrium.⁹³ Just as there is no guarantee that *only* one of them will be accessible, it is an open question whether, say, there really are *either* individualistic or non-individualistic equilibria.

⁹² Perhaps forcing us to accept Lycan’s “double indexical theory of meaning” (“MEANING =_{def} Whatever aspect of linguistic activity happens to interest *me now*.”) In any case, his claim that most accounts of meaning involve “disguised value judgments” (Lycan 1984, 272) would seem, on the model suggested above, to be correct.

⁹³ As is the case with the ‘biological’ conception of fish (assuming that biological kinds are picked out in evolutionary terms).

10. The possibility of no equilibrium points: bifurcationism

This brings us to the unpleasant possibility that there may not only be not many, but, in fact, no equilibria at all. There may be no way of making our general characterizations and putative extensions cohere that does not involve giving up more of either set than we are willing to. We usually suppose that there will be at least *some* coherent way to systematize our linguistic practices (i.e., we not only assume that there will be *at most* one equilibrium point, we assume that there will be *at least* one as well), but this presupposition can turn out to be mistaken.⁹⁴

In such cases there will be a temptation to split the term (as was done with “weight” into “mass” and “weight,” and then “mass” into “rest mass” and “relativistic mass”) into two terms for which equilibria can be reached which together allow us to keep most of our original general characterizations and putative extensions. Those who make a division between ‘narrow’ and ‘wide’ content are effectively accepting such a view.⁹⁵ They conclude that we must split terms like “meaning” and “content” the way we had to split “mass” because they believe that there is no coherent way of systematizing the intuitions and ascriptional practices associated with our semantic terminology. A fairly clear expression of the motivation for such a view can be found in McGinn:

Our intuitive conception of belief content combines two separable components answering to two distinct interests we have in ascriptions of belief. One component consists in mode of representation of things in the world; the other concerns itself with properly semantic relations between such representations and the things represented. I want to suggest that the former component is constitutive of the causal-explanatory role of belief, while the latter is bound up in our taking beliefs as bearers of truth We get different and potentially conflicting standards of individuation -- and hence different conceptions of what a belief essentially is -- according as we concentrate on one or the other component of content. The tendency of discussions of belief is, I think, to allow one component to eclipse the other, thus producing needless conundrums and a distorted picture of the nature of belief.⁹⁶

⁹⁴ For doubts about the ability of linguistic phenomena to accommodate such a presupposition, see, for instance, Austin 1956, 151, Wilson 1982, 552.⁹⁵ See, for instance, Putnam 1981a and McGinn 1982, Loar 1985.

⁹⁶ McGinn 1982, 210-11. While McGinn expresses the motivation for the view very well, his view may not be a case of the view itself, since McGinn’s official line seems to be that while “These components and the concerns they reflect are distinct and independent -- total content supervenes on both taken together.” (210.) Still, it remains the case that, with McGinn, the explanatory work is being done by two independent components, so while there is, technically, just one notion of content for him, there are still two different things explaining behavior and being evaluated as true or false.

McGinn may be a little complacent about how easy it is to give up what he asks us to give up here, but a position such as his might ultimately turn out to be unavoidable. Accepting such “bifurcationism” requires, of course, denying that what it is that is externally determined and evaluated as true or false also plays a part in “commonsense psychological explanation.”⁹⁷ However, while this ‘unity’ thesis has a strong *prima facie* plausibility, it need not be completely nonnegotiable.⁹⁸ It might be something that must be given up since, as long as it remains in our general characterizations, it prevents us from reaching any equilibrium. An equilibrium with it would be preferable to one without, but one without might be preferable to none at all.

It should be noted, however, that not all resolutions involving postulating two notions of content need be ‘bifurcationist’ in a sense that would correspond to a denial of the unity thesis. (Some of these alternatives will be discussed in chapter 2.) The unity constraint is only violated when the two ‘types of content’ split up the explanatory work in the way in which McGinn suggests that they do.

11. Self-reflexive character of dispute when applied to meaning

The dialectic is in many ways more interesting with semantic terminology because the outcome of the dialectic affects how we understand the dialectic itself. After all, the process of bringing the putative extensions and general characterizations for a term like “vegetable” into line does not, in itself, provide a characterization of what is going on in that process. On the other hand, when we become clear about our semantic terms we come to understand the dialectical process as one of either changing what we meant, coming to understand what we always meant, or coming to mean anything at all. The process of bringing the putative extensions and general characterizations of our semantic terms into line will itself provide us with an understanding of what is involved in that very process.

⁹⁷ Bilgrami 1992, 3, 15.

⁹⁸ Though it has occasionally been treated as such. For instance, Bilgrami finds the unity thesis central enough that he is willing to preserve it by “denying not only that meaning is social or normative in any interesting sense, but also by severing all of the standard associations of meaning and content from the official notion of truth-conditions.” (Bilgrami 1992, vii.)

This self-referential quality can be seen by considering two alternative possible equilibria. The first, individualistic, alternative treats what a subject means by a term as constituted by his dispositions to apply the term and the beliefs he associates with it. Someone who thinks whales are fish is understood as simply meaning something different by “fish” than someone who does not (even if this belief does not mesh well with other beliefs of his that also help constitute what he means, such as that “fish” picks out a natural kind and that he means the same thing by “fish” as his compatriots). The second, non-individualistic, alternative treats what the subject means as constituted by whatever the (projected) equilibrium (if any) is for the practices that the society as a whole associates with the term.⁹⁹ Someone who thinks whales are fish simply has a false belief about fish (since that belief is one that would have to be dropped in order to reach an equilibrium).

If the equilibrium for our semantic terms were reached at a comparatively individualistic stage, then the dialectic itself would have to be viewed as a history of the gradual changes in what we meant by our semantic terms (or, at best, the history of how what we meant by our semantic terms became more definite over time). If the equilibrium state were individualistic, then someone who took a resolutely non-individualistic line on semantic terms would have to be understood as meaning something different by their semantic terms than the individualist, and thus could not, by the individualist’s lights, be understood as having false beliefs about the very thing (namely meanings) that the individualist has true beliefs about (though he could, like the person who believed both that whales were fish and that “fish” picked out a natural kind, be criticized for associating an internally inconsistent set of beliefs and practices with his term).¹⁰⁰ On the other hand, if the dialectic ended up at

⁹⁹ These two options are not, of course, meant to be exhaustive. In particular, there could be an individualistic alternative that identified meaning with the equilibrium of the individual’s beliefs and practices, not the beliefs and practices themselves, and there could be a non-individualistic alternative that identified meaning with the socially accepted beliefs and practices, not the equilibrium these beliefs and practices suggested.

¹⁰⁰ So the individualist could not even say the non-individualist claim that Bert meant *arthritis* by “arthritis” was false, since he must understand the non-individualist as meaning something different by “meaning,” “reference” etc.

a very non-individualistic stage, then it might understand itself as the history of how we came to become clear about what we had meant by our semantic terms all along -- the individualist would be understood as simply mistaken about the very thing the non-individualist has come to have a correct conception of.¹⁰¹

Things are more complicated if we come to recognize that there can be more than one equilibrium (say, one for both of the individualist and non-individualist described above). In such a case, the individualist can understand himself as having correct beliefs associated with his semantic terms, even if it is the case the non-individualists beliefs about his semantic terms may be correct as well. Even if everyone else reaches a non-individualistic equilibrium, if the individualist can find an equilibrium for his semantic terms, then he can understand himself as correct. The non-individualist, however, will not only be unable to characterize the individualist as incorrect, but he will also be uncertain whether his own semantic beliefs are correct. Even if there is a non-individualistic equilibrium available for our semantic terms, the non-individualist can only have correct semantic beliefs associated with that equilibrium if that equilibrium is reached by the other members of his community. If everyone else reaches an individualistic equilibrium, then the non-individualist's non-individualism commits him to interpreting himself individualistically.

Finally, there is the case of there being no equilibria at all for our semantic terms, and so, no equilibrium corresponding to either the individualist's or the non-individualist's positions. In such a case, both will have to understand themselves as being, at least to a certain extent, mistaken about their semantic terms. If the lack of equilibria leads to a type of eliminativism, they will both have to understand themselves as very much mistaken, with all of their semantic beliefs being false. If the lack of equilibria leads to a type of bifurcationism, then each may be able to salvage some of their beliefs as true. For instance the bifurcationism might result in a non-individualistic 'wide content' and an individualistic

¹⁰¹ As a result, while advocates of the second way of describing the dialectic can hope to show that their opponents are wrong, advocates of the former can only hope to show that, while their opponents were not wrong about 'the meaning of "meaning,"' they should (because of the place of equilibrium as a regulative ideal) move to a more individualistic conception.

‘narrow content’ with each theorist being able to understand their semantic beliefs as true about one of the two types of content and false about the other.

Discovering, then, how many (if any) equilibria there are for our semantic terms is of considerable importance. We can hope, if not expect, that our intuitions and practices will converge towards a single theory, and finding such a single equilibrium would obviously be of considerable theoretical interest. Finding more, or fewer, than one equilibrium, while in some sense less satisfactory, would be, of course, no less significant. If there are equilibria corresponding to, say, both Davidson’s and Burge’s positions, then debates between the defenders of these two positions remain, in some sense, futile, and what we are faced with are simply instances of indeterminacy. The possibility of there being no or multiple equilibria thus has significant methodological significance for debates about semantic issues, since the fact that there may not be a single equilibrium entails that showing that there is something wrong or ‘incoherent’ about alternative positions is neither a necessary nor a sufficient condition for establishing the adequacy of one’s own position. It is not necessary, since the possibility of there being multiple equilibria allows that one’s position can be adequate in spite of the adequacy of some alternative; it is not sufficient since the possibility of there being no equilibria makes it the case that the ultimate unworkability of all other positions does not entail the acceptability of the final alternative.

12. Conclusion

We have argued, then, that much of the normativity associated with language use arises from the tensions between the putative extensions and general characterizations associated with our terms. Both information-based and inferential role semantics were presented as unsuccessful attempts to explain semantic norms in terms of just one of these two elements that illustrate how the more dynamic relation between the two factors was necessary. It was further argued that a term’s actual extension could be understood as what one gets if one can find a “equilibrium” between the general characterizations associated with a term

and its putative extension. However, when semantic norms are understood this way, a number of problems come to the fore. There is no guarantee that there will an equilibrium to be found for the practices associated with a term, and if one can be found, there is no guarantee that there will be *just* one. Settling, in such cases, whether what a particular term means remains indeterminate until an equilibrium is actually reached for it, or whether reaching an equilibrium makes manifest what the term meant all along turns out to involve getting a proper characterization of our semantic terms. This in turn involves finding an equilibrium for the interplay between their various putative extensions and general characterizations. A number of general beliefs about meaning were discussed, and a number of ascription types brought to our attention by Kripke, Putnam, Burge and Wilson were presented as providing evidence for the putative extensions of these semantic terms. Within this framework, indeterminacy was understood as the product of our being able to find more than one equilibrium for our semantic terms while bifurcationism was understood as the product of our being unable to find any such equilibria. Finally, the way finding an equilibrium for our semantic terms determines how we understand that very process was discussed. The next three chapters will involve discovering just how much of the putative extension made manifest in these ascriptions can be incorporateable into an equilibrium for our semantic terms. It will be suggested that there is an equilibrium that incorporates not only the ascriptional practices brought to our attention by Kripke, Putnam and Burge, but also those discussed by Wilson.