

APPENDICES

Appendix A: Causal theories of reference -- normativity without general
characterizations

It has occasionally been suggested that Kripke and Putnam's 'causal theory' of reference shows how the world's essential structure can help us find a distinction between a term's putative and actual extension without having to resort to the mediation of the general characterizations associated with a word (thus allowing a term's meaning to be determined by its putative extension without being *identified* with it). On such an account, a term's putative extension at the time of its introduction is enough to determine its reference, since what the term refers to can be understood as everything that has the same 'essential structure' as most of the initial sample. A natural kind term picks out the class of things that has the same 'essential' physical structure as most of the items that fall under the term's original putative extension.³²⁹ Proper names, presumably, have the entire person as the initial sample with the extension of the name being determined by the person's essential properties. (That these essential properties are unlikely to map on to any sort of natural kind associated with the biological sciences obviously presents a problem if one is hoping to find a *physicalistic* account of the reference of these terms).

Such an account would have little, if anything, to say about terms that were not natural kind terms, but one could argue that the robust response-independent notion of objectivity we find with purported natural kind terms such as "gold" will not be found with 'social kind terms' such as "chair."

Nevertheless, appealing to causation to give an account of meaning independent of general characterizations is not a very promising project. In particular, for the initial sample 'baptized' to name a person or natural kind, there would seem to need to be some sort of sortal implicit in the baptism (I am not introducing a proper name as a name for the person's family or jaw, and most samples of one natural kind are also instances of others).

³²⁹ For such an account see Devitt 1980 (esp. pp. x, 27-31, 138).

Each baptism makes an at least implicit use of some sortals, and these sortals are reflected in the baptizer's general beliefs about the type of thing that he is baptizing.

It should also be noted that even leaving aside problems relating to the claim that the world has such an 'essential structure,' these purely causal accounts are not especially plausible as readings of either Kripke or Putnam.³³⁰ First of all, the causal reading only really works for natural kind terms (and possibly proper names of people), while it seems fairly clear that the models of language suggested by Kripke and Putnam have wider application. Admittedly, the presupposition that the items picked out in the baptism form a unified 'kind' is stressed most by Kripke with respect to natural kinds,³³¹ but there is nothing in the story presented in *Naming and Necessity* that forces one to say that the only acceptable kinds are natural (rather than functional, legal or culinary) ones. Kripke states that, when the initial "baptism" takes place, "the object may be named by a description,"³³² and while including demonstratives in the description may pave the way for the more indexically driven accounts that purely causal theorists favor, that there be *any* such demonstratives does not seem at all essential to Kripke's picture. (See his discussion of "Neptune.")³³³ Because of this, the reference of terms such as "contract" or "sofa" could be fixed by a description, and speakers unaware of the content of this original description could still refer to whatever the description picked out in virtue of being part of the chain of speakers who used the word. Indeed, Kripke's discussion of the term "meter" seems to presuppose that this is the case. A similar extension to non-natural kind cases can be found in Putnam's proposals about the 'hidden indexicality' in our language.³³⁴ Admittedly, Putnam will talk about "hidden structures,"³³⁵ and tends to focus on natural kind terms

³³⁰ A claim defended recently, and with semantic issues in mind, by David Lewis. (See his discussion of 'elite properties' in Lewis 1983.)

³³¹ See, for instance, Kripke 1972, 120-1, 134-6, 138.

³³² Kripke 1972, 96.

³³³ Kripke 1972, 79.

³³⁴ Putnam 1975, 234, 271.

³³⁵ Putnam 1975, 241.

like “gold,”³³⁶ but he goes on to claim not only that “‘pencil’ is just as indexical as ‘water’ or ‘gold,’”³³⁷ but also that his conclusions “apply to the great majority of all nouns, and to other parts of speech as well.”³³⁸

Furthermore, it seems clear that, even in the case of natural kind terms, the reference of the terms in question do not correspond to the natural kind instantiated by most of the items that fall within the putative extension because the way science carves up the world is ‘metaphysically privileged’ in a way that allows the world’s ‘essential’ structure to impose itself on our words. Rather, it does so because the type of structures discovered by science correspond to the deeply entrenched presuppositions we have about the types of things we are talking about when we use kind terms. That is to say, it is our general characterizations that help determine which parts of the putative extension are ‘really’ correct. As a result, if we had a different set of general characterizations, different types of kinds would be picked out by exactly the same mechanisms that picked out the natural ones. For instance, Putnam claims that the ‘hidden structures’ determine the reference of natural kind terms, not because only such hidden structures could enter into causal relations with the baptizer, but rather because “normally the ‘important’ properties of a liquid or a solid, etc., are the ones that are *structurally* important.”³³⁹ The hidden structure of the initial sample determines the extension of the term because hidden structures are what are important to the baptizer and the speakers of the language. This view has as an obvious consequence that if other types of features are taken to be more important by the baptizer (and, as Putnam stresses, “importance is an interest relative notion”),³⁴⁰ other types of features would be treated as ‘essential’ and thus determine how the initial set of examples is to be extended to a kind.

This difference of interest can be found even among those who agree that the term should be treated as a natural kind term. Consider, for instance, Putnam’s discussion of

³³⁶ Putnam 1975, 235.

³³⁷ Putnam 1975, 243.

³³⁸ Putnam 1975, 242.

³³⁹ Putnam 1975, 239.

³⁴⁰ Putnam 1975, 239.

how a molecular biologist and an evolutionary biologist might disagree over what was ‘essential’ to being a dog. A ‘dog’ that was found to be molecularly and genetically indistinguishable from ‘real’ dogs, but was known to have a completely separate evolutionary history, would be considered a dog by the molecular biologist’s conception of species identity, but not by the evolutionary biologist.³⁴¹ Non-natural kind terms can thus all be accounted for in the baptismal picture simply by allowing that a baptizer can have different interests, and thus consider different types of properties ‘important,’ when they introduce new terms. These different interests can, of course, be understood in terms of (possibly implicit) general characterizations associated with the terms, and so if interests are brought in, one ends up with a dynamic picture.

That something like general characterizations are in the background of Kripke’s picture is also clear when he says that “even though we don’t *know* the internal structure of tigers, we suppose... that tigers form a certain species or natural kind.”³⁴² It is this supposition that determines what parts of the putative extension do and do not fall under the term’s actual extension. A similar thought is found in his suggestion that everyone could be wrong about whales being fish because responsibility to the internal structure of the animals in question was “part of the original enterprise” of classification.³⁴³ Furthermore, even if one was willing to admit that scientific taxonomies were somehow ‘more real’ than non-scientific ones, this itself would do nothing to show that our everyday terminology ‘really’ corresponds to some scientific taxonomy.

Kripke’s and Putnam’s responses to the overemphasis on general characterizations found in theories like the descriptive theories of names is thus best understood as providing

³⁴¹ Putnam 1993, 77.

³⁴² Kripke 1972, 120-1. Furthermore, both Kripke and Putnam are willing to say that these general characterizations which structure our practice can turn out to be mistaken. Kripke is happy to admit that “we may be mistaken in supposing that there is such a kind” (121), and when we have made such a mistake, we would not go on to say that there are no tigers. The same line of thought can be seen in Putnam’s suggestion that, in spite of the fact that we assume that they are animals, we could imagine discovering that all of them are robots.

³⁴³ Kripke 1972, 138.

a more dynamic picture of what was involved, not as advocating a theory of meaning where the putative extensions (with the help of the world's 'essential' structure) are overemphasized in a corresponding fashion. Given that they were reacting to one extreme (meaning determined solely by general characterizations), it is not surprising that their own work can be easily (mis)understood as advocating a recoil to the corresponding extreme (general characterizations play no role at all), but such a reading not only makes their views seem fairly implausible, but also reduces their significance by missing the extent to which the stories about natural kinds and proper names can be understood as particular examples of a much more general phenomenon.

Appendix B: Semantic defects as epistemological virtue

Accounts that tie meaning too closely to putative extensions while ignoring the general characterizations are unable to give any substantive account of the possibility of our being wrong about the correctness of our application to those items in (the core of) the putative extensions. However, there has been a tradition of trying to make a virtue out of this defect in one's semantic theory. In particular, if one's account cannot make sense of the possibility of one's being mistaken in a certain range of cases, then the account leaves no room for skeptical doubts about these cases. As a result, rather than being criticized for underfunding our conception of objectivity, such accounts of meaning have on occasion actually been praised for their purported 'anti-skeptical' consequences.

Tying what a term means to its putative extension³⁴⁴ and then appealing to what the term means to fend off skepticism is characteristic of so-called 'paradigm case' arguments. Paradigm cases tend to be the most deeply entrenched members of a term's putative extension, so if what a term means is solely a function of its putative extension, there will be no way for these central applications to turn out to be mistaken. An instance of this strategy can be seen in Stebbing's response³⁴⁵ to Eddington's claim³⁴⁶ that desks were not 'really' solid.

What are we to understand by "solidity"? Unless we do understand it we cannot understand what the denial of solidity to the plank amounts to. But we can understand "solidity" only if we can truly say that the plank is solid. For "solid" is just the word we use to describe a certain respect in which a plank of wood resembles a block of marble, a piece of paper, and a cricket ball, and in which each of these differs from a sponge, from the interior of a soap-bubble, and from the holes in a net. ³⁴⁷

³⁴⁴ Or at least the 'core' of the putative extension: these would be the parts of the putative extension that are the most deeply 'entrenched,' perhaps the applications made in certain 'optimal' conditions (for something like "that is red") or items that are 'prototypical.' Of course, such deeply entrenched items will be much less likely to be given up than others, but being wrong in such cases is not impossible. If they are given up, then perhaps the rest of the putative extension will be given up too (i.e.: nothing will be taken to satisfy the term in question). Still, while one will expect this to be a fairly rare occurrence, it can still happen (even 'paradigm cases' of witches will turn out not to be, since there are, of course, no witches).

³⁴⁵ Which is endorsed by Urmson as a successful application of "the argument from standard examples." (Urmson 1953, 121-2.)

³⁴⁶ In Eddington 1929.

³⁴⁷ Stebbing 1937, 51-2.

Even if one does agree that physics has not shown that tables are not solid, the argument from paradigm cases is far from conclusive. It relies on the meaning of terms like “solid” being determined solely by their putative extensions, and so leaves itself open to responses such as the following:

The effectiveness of Eddington’s remark seems to derive from the fact that the properties which one is inclined to regard as the defining properties of solidity do not really belong to the objects to which we customarily apply the term ‘solid,’ and the interest of this fact is in no way diminished by repeating ‘Oh, but we *do* apply the word “solid” to such things as desks.’ Of course we do, but we also say that by ‘solid’ we mean ‘having its interior entirely filled with matter,’ or something like that. The collision is *within* ordinary usage, and not between it and scientific theory.³⁴⁸

Eddington took physics to have shown that most of what fell under the putative extension of ‘solid’ did not fit the general characterization we associated with that term. While Eddington may have given too much to the general characterization of “solid,” Stebbing focused too much on the term’s putative extension. The question of whether tables are *really* solid or not will inevitably involve comparing both.

And if the two cannot be reconciled, (i.e.: there is no real equilibrium involving both) one might want to say that there is, in fact, no equilibrium for our previous use of “solid,” and that the term needs to be replaced by two homonyms, one tied to the putative extension, another to the general characterization. Both the claim that science has shown that tables aren’t ‘really’ solid and the claim that it has shown that solid objects ‘really’ consist mainly of empty space presuppose that an equilibrium can be found that dumps most of either the general characterizations or the putative extension.

Much recent work in epistemology informed by the ‘externalist’ semantic tradition also seems to focus too much on the putative extensions of our terms, particularly with reference to the central and ‘paradigmatic’ examples.³⁴⁹ Externalism about content

³⁴⁸ Mates 1958, 127. Mates distinguishes two approaches to discovering what a person means by a given term. The “extensional” approach (which focuses just on the putative extension the speaker associates with the term) and the “intensional” approach (which attempts to get the speaker to come up with general characterizations). I am fairly sympathetic not only with this distinction, and but also with his suggestion that the results of these two methods need not coincide. (See Mates 1958, 125-6.)

³⁴⁹ Similar epistemic work can be done, of course, by overemphasizing the aspects of the general characterizations. Take, for instance, the following example from C.I. Lewis:

encourages people to think that systematic perceptual error is impossible because the contents of our perceptual terms are determined precisely by those objects we apply those perceptual terms to.³⁵⁰ Nevertheless, while it may be true that widespread *perceptual* error could be ruled out, this certainly does not mean that the perceptually-based judgments must be true. There have been many witch sightings, and, since there are no witches, we can assume that all these perceptually based witch-beliefs were false. What is at issue here is not ‘perceptual error,’ rather it is that our general characterization of “witch” was not satisfied by any of the people who were called “witches.” That the meanings of our terms could be affected by these general characterizations makes the transition from the impossibility of general perceptual error to anti-skepticism about perceptual judgments problematic.³⁵¹

Consider, for instance, a speaker who refers, as we do, to both the morning star and the evening star as “Venus.” When Venus rises at the beginning of the evening he may say “there is Venus.” When he is asked what he is looking at, he will reply “Venus.” He does the same in the morning, and never ‘misapplies’ the term to any other star or planet. The putative extension of “Venus” in his language consists only of the planet Venus. However, while the putative extension of Venus in his language is much like it is in ours (assuming we can identify Venus in the night sky), the general characterization he associates with the term is very different. In particular, while we believe that Venus is a lifeless planet, that we have sent satellites to explore it, etc., this speaker believes that Venus is the goddess of love, that she is married to Vulcan, is the lover of Mars, responds to the prayers of her devoted followers (of whom the speaker is one), etc.

It is an a priori principle that physical things have mass. By this criterion they are distinguished from mirror-images and illusion.... the truth of the principle is independent of the particular phenomenon. But a world in which we should experience phenomena having persistence and independence not characteristic of imagination, and a coherence not characteristic of our dreams, but things which would still not be amenable to any gravitational generalizations, is entirely conceivable. In such a world our a priori principle would not be rendered false -- since it is definitive of the physical; but the category “physical” might well be useless. (Lewis 1929, 26-7.)

³⁵⁰ See, for instance, Burge 1986b, 131.

³⁵¹ This is, I take it, the gist of Nagel’s complaint about Putnam’s Brain-in-the-vat argument. (Nagel 1986, 71-3.)

Now it should be fairly clear that for his term “Venus,” there is a serious lack of fit between the general characterizations and what he actually applies the term to. If the speaker is a particularly devout worshiper, he may be much more willing to give up everything in the putative extension rather than his general beliefs. If he finds out that ‘Venus’ is a lifeless planet, he will be inclined, correctly, to think that he misapplied the term every morning and evening when he looked at the sky.³⁵²

The holistic nature of the dynamic between our terms’ putative extensions and general characterizations may insure that, if our terms pick out anything at all, the most deeply entrenched elements of our usage will turn out to be true, but there is no guarantee that anything from a particular term’s putative extension will be among the most deeply entrenched set.³⁵³

³⁵² There may be some temptation to treat utterances of his such as “Venus is looking especially bright tonight” as still being true (provided that the planet is especially bright on that night), but this is because it is true in the sense that my utterance of “John looks hurt” might be considered true if I mistake Peter for John and Peter looks hurt. The speaker’s reference for such utterances may be Peter or the planet, but the semantic reference will still be John or the Greco-Roman goddess.

³⁵³ Nor, for that matter, will there be any guarantee that anything from the general characterization will be among such a set.

Appendix C: On “The Methodology of Naturalistic Semantics”

In his recent paper,³⁵⁴ Devitt points out that “we start semantics in the unusual position of having to specify a subject matter,” and while he thinks that “we should not insist on great precision about this in advance of theory,” he still insists that “we do need some explication of our vague talk of ‘meanings’ ... that is *not ad hoc*.”³⁵⁵ He conceives of this as just one instance of a more general type of problem:

What constitutes some property we ascribe, “being an F”? The question may often have little scientific interest, or course, but where it does -- presumably “being a gene,” “an atom,” “an acid,” “an echidna,” and “a pain” are examples -- *how do we tell*? Sometimes we already have a well-established theory; for example, thanks to molecular genetics, we have for “being a gene.” But suppose that we do not and are starting pretty much from scratch, which is surely the right supposition for semantics. How then do we tell what is common and peculiar to Fs “in all possible worlds?” What is the “ultimate” method?³⁵⁶

Devitt describes this “ultimate” method as follows:

It breaks into two stages. First we must identify some apparently uncontroversial examples of F’s and non-F’s. Second, we must examine the examples to determine the nature of “being an F.” ... This second stage is a straightforwardly scientific one. It is hard to accomplish and, as is well known, even harder to say how we accomplish it We should note, however, an important feature of this second stage: the examination can lead us to reject some of the results of the first stage; apparently uncontroversial examples turn out to be controversial. Thus we may conclude that some of the things identified as Fs are not; for example, whales are not fish. We may conclude that some things identified as non-Fs are Fs; for example, tomatoes are fruits. We may even conclude that *none* of the things identified as Fs is an F; for example, there are no witches. Usually, such revisions in our preliminary identifications take place only if our examination leads to a powerful theory.³⁵⁷

While the contents of this chapter commit me to treating something like this proposal of Devitt’s as right, Devitt’s own characterization of the ‘ultimate method’ has a number of serious flaws that prejudice his account of semantics (or perhaps simply reflect antecedent prejudices about semantics).

As the title of his paper suggests, Devitt wants his account of semantics to be “naturalistic,” and is unapologetic about wanting to “bring semantic methodology close to other scientific methodologies,” since, after all, “according to naturalism, semantics is an

³⁵⁴ Devitt 1994.

³⁵⁵ Devitt 1994, 548.

³⁵⁶ Devitt 1994, 561-2.

³⁵⁷ Devitt 1994, 562.

empirical science like any other.”³⁵⁸ This naturalistic methodology is supposed to, among other things, constrain our reliance on intuitions in semantics:

Clearly, there could be no objection to the standard appeal to intuition in semantics if the semantic task were simply the systematization or explanations of the intuitions. Some write as if they think that this is the task. Yet this task should be no more interesting or appropriate than that of systematizing or explaining our ordinary physical, biological, or psychological intuitions. The intuitions in each case are *about* the subject matter that should concern the science, they are not the subject matter themselves.³⁵⁹

In the face of the dominant role played by intuitions in semantics, it is worth laboring the point. From the naturalistic perspective, semantic intuitions are like those in any other science: open to revision in the face of empirical theory. We could be wrong about what has putative meaning. We could be wrong in thinking that anything has it. Even if we are right that something has it, we could be wrong in thinking that it plays a semantic role and so really is meaning.³⁶⁰

Intuitions associated with most of the ‘standard’ thought experiments thus play a role in the first stage of the “ultimate method,” since these attitude ascriptions will be the source of the “uncontroversial examples” to be gathered during the first stage.³⁶¹

Devitt is right to think that our semantic intuitions associated with various thought experiments should be understood as playing something like this role, and his complaints about the *ad hocness* of various semantic proposals that do not follow anything like the ultimate method seem justified. Nevertheless, his own insistence that semantics be naturalistic seems to be an instance of precisely the type of tenacious devotion to an *a priori* claim about content that an ultimate method should free one from. His version of the ultimate method cannot do this since it is only supposed to work for natural kind terms.³⁶² This becomes especially clear when we consider that the second stage is supposed to be “a straightforwardly scientific one.”³⁶³ As a result, the method would only work for natural

³⁵⁸ Devitt 1994, 545.

³⁵⁹ Devitt 1994, 546. See also, “the ‘basic’ task is to say what meanings *are*, to explain their *natures*.” (Devitt 1994, 547.)

³⁶⁰ Devitt 1994, 568. “Putative meanings” are discovered by the “descriptive task” of giving an account of how we use our semantic terminology. Deciding if these putative meanings play the semantic roles of explaining behavior and providing truths about the world (Devitt 1994, 550) (and thus if they are *real* meanings) is taken to be part of the “normative task” (Devitt 1994, 565).

³⁶¹ Devitt 1994, 566.

³⁶² “Gene,” “atom,” “acid,” “echidna,” “pain,” are all, presumably, taken to be such terms.

³⁶³ Devitt 1994, 562.

kinds, and thus, if it is to work for semantic terms at all, semantics would have to be naturalistic.

However, since the assumption that semantic terms should be understood naturalistically is undermotivated, the fact that Devitt's version of the ultimate method requires this assumption suggests that Devitt's characterization of the ultimate method is not general enough.³⁶⁴ Indeed, we can find a more general account of how we find "what constitutes some property we ascribe, 'being an F'" in Tyler Burge's "Intellectual Norms and the Foundations of Mind,"³⁶⁵ with Burge's "examples" and "archetypical applications"³⁶⁶ corresponding to Devitt's first stage and his "normative characterizations" corresponding to something like Devitt's second stage. For a natural kind term, the task of coming up with a normative characterization based on the examples may, of course, be a "straightforwardly scientific one," but Burge makes it clear that the method is intended to apply equally well to such non-natural kinds as "chair," "baby," "walk," or "knife."³⁶⁷

In addition (and related) to its lack of generality, Devitt's account of the ultimate method suffers from treating the dialectic between the uncontroversial examples and general characterizations too simply. In particular, it treats the dialectic as if the general characterizations came only after scientifically investigating (or at least reflecting upon) the uncontroversial examples. This seems unrealistic; for most terms we will have both a set of uncontroversial examples of and uncontroversial beliefs about the items in question, and any ultimate method will involve finding an equilibrium between these two. Take, for instance, Devitt's own example of "witch." It does not seem possible that we came to the conclusion that there were no witches by anything like the method Devitt proposes. In particular, it seems highly unlikely that we first identified a number of 'uncontroversial

³⁶⁴ Devitt seems, at times, quite happy to admit the arbitrariness of this assumption (Devitt 1994, 558).

³⁶⁵ Burge 1986c.

³⁶⁶ Burge 1986c, 703.

³⁶⁷ And, of course, Rawls can be understood as applying something like the ultimate method for ethical terms like "justice" ("reflective equilibrium" being the ideal point where the products of the two stages correspond (Rawls 1971)). Indeed, this general 'Socratic' procedure has a fairly long and distinguished history.

examples of witches' and then through scientifically investigating these examples came up with a characterization of what witches are that none of these examples satisfied. Furthermore, even if there were a property that all and only the purported witches had (red hair, warts on the nose, non-matching eyes, whatever), we would still have concluded that there were no witches, not that "witch" turned out to be synonymous with, say, "redhead." We concluded that there were no witches because we have a number of strong general beliefs about witches (they can cast spells, work for the devil, etc.) that we are not inclined to give up even if none of the "uncontroversial examples" of witches turn out to satisfy them.

There can be a number of such general beliefs that are brought into an investigation rather than derived from it, and the assumption that a term picks out a natural kind may be just such an assumption.³⁶⁸ Furthermore, just as reflection and general theories about what it is to be an F can cause us to reject some of the originally uncontroversial examples, they can also cause us to reject some originally uncontroversial assumptions. The assumption that a term picks out a natural kind is just such an assumption, and there are many terms (for instance, "tree," "lily," "reptile," and "fish") for which this assumption has turned out to be false.

When such presuppositions conflict with the results of investigation, we may have a number of options. Consider the case of the word "ape." This term was applied uncontroversially to oranges, chimps and gorillas. It was also not applied to humans (we were descended from apes, but we were not apes ourselves). The term was also assumed to pick out a natural kind. Unfortunately, there is no biological kind that includes oranges, chimps and gorillas and does not include humans as well.³⁶⁹ Given this fact, we are faced with three options: (1) admitting that we are, in fact, apes (treating "apes" being a natural

³⁶⁸ That the 'real' semantic properties explain behavior and that they tell us about the world are also presupposed by the investigation rather than derived from it, though one would expect that these general beliefs are a good deal less negotiable than Devitt's naturalism.

³⁶⁹ Gould 1983, 358.

kind as non-negotiable, our not being apes as negotiable), (2) admitting that “apes” does not, after all, pick out a natural kind (treating “apes” being a natural kind as negotiable, our not being apes as non-negotiable), or (3) admitting that there are no apes (treating both as non-negotiable). By building naturalism into the ultimate method itself, Devitt would force us to take either the first or the third option, even though the second is, in many ways, the most appealing.

If the idea that semantic terms pick out natural kinds (i.e.: that semantics can be “naturalistic”) is just one general idea that can be brought into the dialectic associated with any real ultimate method, then it obviously should not be a precondition of any ultimate method being used at all. Though one could, of course, take the idea that semantic terms pick out natural kinds to be non-negotiable in the way in which witches’ casting spells may have been,³⁷⁰ in which case, it could simply turn out that there were no meanings, beliefs, etc.³⁷¹ This option, which could push one towards a form of eliminative materialism, does not seem especially desirable, and while it is just an unappealing option on the more general method described above, it seems as if it could very well turn out to be a necessity on Devitt’s. A more general methodology would make ‘folk psychological’ terms like “belief”

³⁷⁰ Though reasons for thinking that we should not treat naturalism as being non-negotiable are given by Devitt himself:

Suppose that our actual practice of ascribing certain properties serves our purposes in some area well. That is very good evidence that we ought to ascribe at least those properties to serve those purposes. The situation in semantics appears to be of this sort. For consider how successful our ordinary attitude ascriptions seem to be at serving their purposes. Their apparent success in explaining behavior is certainly not limitless but it is nonetheless impressive. Their apparent success in influencing our own thoughts about the world could hardly be exaggerated. So it is likely that, at least, we ought to ascribe to tokens that are thought and uttered the properties that we do ascribe and hence that those properties are meanings. Evidence for a descriptive theory will be evidence for a normative one and hence the basic one In the absence of a detailed comparison with the semantic status quo, a revisionist proposal will have a credibility problem: it will always seem much more plausible that there is a flaw in the philosophical arguments offered in support of the proposal than that the apparently successful status quo is so radically mistaken. (Devitt 1994, 561)

The assumption that an account must be naturalistic may be just the type of flaw that exists in arguments attempting to overthrow the status quo. After all, the success of a practice certainly doesn’t imply that the terms of the practice will necessarily map on to natural kinds (functional kinds being just one obvious alternative around which successful practices could be based).

³⁷¹ An outcome that Devitt seems to consider extremely unlikely, though still possible (Devitt 1994, 561).

no more threatened by the possibility of a failure of naturalistic reduction than 'folk biological' terms like "tree" or "lily."

Appendix D: Methodological and Ascriptional Individualism.

The purpose of this appendix is to distinguish “methodological” individualism from “ascriptional” individualism. In the body of the chapter, the term “individualism” is used to pick out ascriptional individualism, and while this sort of individualism is rejected, methodological individualism is not criticized. It will be argued in this appendix that writers such as Davidson and Bilgrami occasionally defend ascriptional individualism because they take their commitment to methodological individualism to entail it.

Methodological individualism requires that factors ‘external’ to the interpretee (such as, say, the usage of the surrounding community or the physical structure of what is around him) be counted as relevant in determining what he means only if such factors are taken to be relevant by the interpretee (in its crudest form, he takes “arthritis” to mean “what the experts mean by ‘arthritis,’” or he takes “water” to mean “whatever has the same atomic structure as most of the stuff I’ve called ‘water,’” etc.). Methodological individualism is thus incompatible with a strong anti-individualistic position according to which we can be assured, *a priori*, that a speaker’s words should be interpreted according to the community standard regardless of the speaker’s attitude towards such a standard. In Davidson’s terms, methodological individualism requires that all linguistic norms must ultimately be grounded in the intentions of the agent himself.³⁷² Davidson puts forth a methodologically individualist position in the following passage.

If the speech behavior of others does not provide the norm for the speaker, what can? The answer is the intention of the speaker to be interpreted in a certain way provides the norm; the speaker falls short of this intention if he fails to speak in such a way as to be understood as he intended. Under usual circumstances a speaker knows he is most apt to be understood if he speaks as his listeners would, and so intends to speak as he thinks they would. He will then fail in one of his intentions if he does not

³⁷² A charitable reading of *Naming and Necessity* is thus compatible with methodological individualism. While Kripke is, justly, unsympathetic with metalinguistic attempts to give a sense to proper names, as stated earlier, it is an *essential* part of his story that “the speaker intends to use the name in the same way as it was transmitted to him” (i.e., he intends to be a ‘reference preserving link’ in the chain of communication). There are problems with Kripke’s initial formulation of the idea (see appendix F), but it is clear that, for Kripke, the relevant norms only get a grip on the speaker through the mediation of his intentions. This was, of course, an aspect of Kripke’s account that Devitt tried to get rid of in giving his purely ‘causal’ account. As a result Kripke’s position, unlike Devitt’s, could still be understood as methodologically individualistic. (Devitt’s cannot because he suggests that the ‘real essences’ of the kinds in our environment affect what we mean independently of any concern we have with them.)

speak as others do. This simple fact helps explain, I think, why many philosophers have tied the meaning of a speaker's utterance to what others mean by the same words ... (where "others" refers to a linguistic community, "experts," or an élite of one sort or another). On my account, this tie is neither essential nor direct; it comes into play only when the speaker intends to be interpreted as (certain) others would be. When this intention is absent, the correct understanding of a speaker is unaffected by usage beyond the intended reach of his voice.³⁷³

This claim about the indirectness of the connection is quite plausible. While we are usually willing to consider ourselves subject to the norms of the surrounding community, if we decide, in certain cases, deliberately to flout communal usage it is hard to see why what we mean should still be tied to communal usage. If I insist on calling stools "chairs" in spite of having the fact that my use is non-standard pointed out to me on many occasions,³⁷⁴ and I make clear to people unfamiliar with my usage how it differs from the norm,³⁷⁵ then it seems perverse to describe my calling a stool a "chair" as the expression of a false belief. It would be better to say that I just meant something different by "chair" than the rest of my community.

Still, one can agree with Davidson that if one lacks the intention to be interpreted as others would be, then one shouldn't be so interpreted, and still argue that in most cases, we do intend to be interpreted as others would be.³⁷⁶ Claiming that the norm to which I am

³⁷³ Davidson 1992, 261.

³⁷⁴ I reply in these cases "Well, I just mean something different by "chair" than most people do."

³⁷⁵ Davidson stresses how important it is that I make sufficient effort to make my idiosyncratic usage known.

If a speaker reasonably believes he will be interpreted in a certain way, and speaks with the intention of being so understood, we may choose to say he means what (in the primary sense) he would have meant if he had been understood as he expected and intended. Reasonable belief is such a flexible concept that we may want to add that there must be people who would understand the speaker as he intends, and the speaker reasonably believes he is speaking to such a person.... These remarks should make plain why Dummett's accusation, that I endorse a variety of Humpty Dumpty's theory that meaning depends only on intention, does not find its target. (Davidson 1994a, 12-13.)

Still, if one really must be able reasonably to expect others to be able to interpret one correctly, one might think that departures from the social norm could only be exceptions made intelligible against a background of general conformity, and that one could not reasonably expect to be properly understood if one departed from the surrounding norms for *all* (or even most) of one's words. Especially if, like Davidson, one thinks that disagreement is only intelligible against a shared background of agreement.

³⁷⁶ This need not, of course be an explicitly formulated prior intention, but can be seen as part of a background to our linguistic intentions (see appendix F), which leads us to understand ourselves as mistaken when we defer to accepted usage. (Indeed, deference itself may be seen as involving this intention as a type of 'intention in action.' Provided, of course, one thinks the concept of an intention in action makes sense -- which Davidson clearly does.)

beholden in using a word is determined by how I intend to use the word³⁷⁷ is not incompatible with non-individualism if how I intend to use the word usually turns out to be tied to the linguistic practices of the surrounding community.³⁷⁸ In particular, we can assume that Bert intended to be interpreted as the doctor would be, and his failed intention to speak correctly did foil “the intention to be interpreted a certain way.” Indeed, the fact that Bert is speaking to a *doctor* and not, say, his brother is relevant here. When we are actually speaking to experts in a context in which their expertise is being appealed to, the presumption that we will intend to be following ‘expert usage’ will be considerably stronger than when we are talking within a possibly idiosyncratic sub-community such as our family. There is little reason to think that someone couldn’t accept Davidson’s point and still insist that *Bert* should be interpreted the way Burge does.³⁷⁹ It may be the case that, when Bert complains to his brother, “I hate this cold weather, it always makes my arthritis worse,” he should be interpreted as meaning *tharthrititis* by “arthritis,” but the case Burge describes is not like this at all. Bert is going to see a *doctor* and in this context (consulting a doctor) it seems quite plausible to think that Bert intends to be following expert usage, his utterance of “I have arthritis in my thigh” having (very roughly) the gist of “I know there is this disease you doctors call “arthritis” and I think that I have *it* in my thigh.” Such a ‘metalinguistic’ reworking of his communicative intention has little plausibility in the case of his complaint to his brother, and Burge’s choice of context has a lot to do with the plausibility of his example. After all, if he *didn’t* believe that the doctor knew more about arthritis than he did, it isn’t clear why he would be going in for a

³⁷⁷ Intentions “give content to an attribution of error by allowing for the possibility of a discrepancy between intention and accomplishment.” (Davidson 1992, 259.) The appeal to community is, of course, not the only way to generate cleavage between how I intend to use a word and how I actually use it. Any criteria by which I can ‘take back’ various attributions of the term as ‘mistaken’ will be enough to generate the difference.

³⁷⁸ Of course Davidson claims that “a failed intention to speak ‘correctly,’ unless it foils the intention to be interpreted in a certain way, does not matter to what the speaker means” (Davidson 1992, 261), but the non-individualist can respond that “our intention to be interpreted in a certain way” is often directly connected to our intention to speak ‘correctly.’

³⁷⁹ After all, Burge himself claims that “the individual’s intentions or attitudes toward communal conventions and communal conceptions seem more important than the causal antecedents of his transactions with a word” (Burge 1979, 114).

consultation in the first place. The tie between the meaning of a speaker's utterance and what others mean by the same words may be neither essential nor direct, but it is often there nonetheless.

Methodological individualism is thus compatible with our occasionally ascribing contents in a fashion that this chapter characterizes as "non-individualistic." We can say, for instance, that Bert means *arthritis* by "arthritis" if he intends to use the term to pick out what the doctors do by it. Davidson endorses such a possibility in passages such as the following:

I do not doubt the existence of [the linguistic division of labor], or even its importance....we can take it to be part of the meaning of an expression that its reference is to be determined by expert opinion.... So for the words "elm" and "beech" to pick out the appropriate trees there would have to be experts, but we cannot conclude that the meaningful use of these words demands a social setting.... It is obvious that the linguistic division of labor can't be essential to verbal communication. If everyone meant by "elm" "what others mean by elm," the word would have no reference. The linguistic division of labor is a device that can come into play only after the basic linguistic skills that tie words directly to things are already in place. So no matter how universal the linguistic division of labor is in practice, it cannot constitute the essential social element in language. We could get along without it.³⁸⁰

Davidson here admits that, for *most* of us, what we refer to by terms like "elm" may be determined by experts,³⁸¹ so if the experts had referred to something different, we would have too. These 'non-individualistic' content ascriptions are, however, justified on purely *individualistic* grounds. 'Accepted' usage is relevant only because the *individual* takes it to be. Davidson seems only to be denying that, for any given word, this can be true of *all* of us, and thus that the division of linguistic labor is essential to language as such. His real target is, then, not the claim that the truth conditions of our thoughts and utterances may be in part determined by the usage of the surrounding community, but rather the stronger claim that such dependence is a *necessary* feature of language use.³⁸²

³⁸⁰ Davidson 1994a, 5-6. In much the same way, Bilgrami seems to endorse the claim that we can ascribe contents non-individualistically if the speaker himself intends to be interpreted that way. (See Bilgrami 1992a 23, 42.)³⁸¹ Davidson, of course, contests this very point in a number of places, including the pages immediately following the passage quoted above.

³⁸² And this claim about language and communication *as such*, need not be in conflict with Burge's claim that the fact that we "normally look to others to help set standards for determining the range of legitimate examples and the sort of background information used in explicating a word or concept derives from psychological necessities for human beings." (Burge 1989, 186, 187).

These certainly aren't the only places where Davidson seems to commit himself to a position that is ascriptionally non-individualistic. Davidson not only claims that "a social theory of interpretation" is made possible by belief taking up "the slack between sentences held true by individuals and sentences held true (or false) by public standards,"³⁸³ but also argues that "states of mind like doubts, wishes, beliefs, and desires are identified in part by the social and historical context in which they are acquired."³⁸⁴ Furthermore, ascriptional individualism does not sit well with Davidson's discussion of "porcupine" in the following passage:

I may not know the difference between an echidna and a porcupine; as a result I may call all the echidnas I come across porcupines. Yet, because of the environment in which I learned the word 'porcupine,' my word 'porcupine' refers to porcupines and not echidnas; this is what I think it refers to, and what I believe I see before me when I honestly affirm, "That's a porcupine."³⁸⁵

It is not at all clear what the relevant structural difference is supposed to be between Davidson's porcupine/porcupine-or-echidna example and Burge's arthritis/tharthritis example.

Indeed, Davidson here seems to commit himself to there being cases where what we are talking about should be understood determined by our social environment. Davidson's reference to "the environment in which I learned the word" should not be understood as a reference merely to his *physical* environment, since Davidson cannot be suggesting that he refers to porcupines because he learnt the word "porcupine" when he was in contact with members of the natural kind *porcupine*.³⁸⁶ Not only has he expressed skepticism about the idea of natural kinds,³⁸⁷ but he also explicitly insists that the phenomenon is not limited to natural kind terms and is "ubiquitous, since it is inseparable from the social character of language."³⁸⁸

³⁸³ Davidson 1974a, 153. He further claims that belief is a private attitude that "is not intelligible except as an adjustment to the public norm provided by language" (Davidson 1975, 170).

³⁸⁴ Davidson 1989, 170.

³⁸⁵ Davidson 1989, 168.

³⁸⁶ Which, for most of us, would be an implausible claim about how the term was learnt.

³⁸⁷ Davidson 1987, 450.

³⁸⁸ Davidson 1989, 167.

Ascriptional individualism leaves no room for the intermediate position that methodological individualism allows for. In particular, if Bert believes that he has “arthritis” in his thigh and applies the term to a pain in his thigh, then the ascriptional individualist will claim that there is no way to understand him as having misapplied his term, and thus no way to understand him as referring to *arthritis* by “arthritis.” On such an account, it is not merely the case that semantic (lexical) normativity must ultimately be grounded in the speaker’s intentions, but also the case that there is no room for semantic norms to really go beyond perceptual errors. Ascriptional individualism is thus considerably stronger than methodological individualism, and it is incompatible with the type of ascriptionally non-individualistic position that the methodological individualism allows for.

Both Bilgrami and Davidson, however, occasionally advocate versions of ascriptional individualism in passages such as these:

I am not impressed by [Dummett’s] or Burge’s or Putnam’s insistence that words may have a meaning of which both speaker and hearer are ignorant. I don’t doubt that we sometimes say this, and it’s fairly clear what we have in mind: speaker and hearer are ignorant of what would be found in some dictionary, or how people with a better or different education or a higher income use the words. This is still meaning based on successful communication, but it imports into the theory of meaning an elitist norm by implying that people not in the right social swim don’t really know what they mean.³⁸⁹

The notion of a mistake is a normative notion, and the possibility of mistakes of this sort introduces a matter of rightness or wrongness regarding the use of terms or the application of concepts the terms express. Rightness or wrongness of this sort is essential to the idea of meaning and concepts, and therefore to the idea of content. My view does not allow one to say of Bert and others like him that they have made or can make these sorts of mistakes.... when it came to describing Bert’s deference to his doctor, I did not describe it as learning from the doctor that he was all along mistaken about arthritis and had been wrongly applying his term for it, but rather as the doctor having helped him acquire a new concept so that he was now wielding a new term. Therefore my view has no place for an essential normativity that attaches to meaning and content.³⁹⁰

While Davidson and Bilgrami defend both ascriptional and methodological individualism, their motivation seems to be tied more to methodological individualism, and both endorse claims that are incompatible with ascriptional, but not methodological, individualism. Indeed, the considered positions of both of them are probably versions of methodological

³⁸⁹ Davidson, 1994a, 6, 9. It should also be noted that, in those cases where we use a word idiosyncratically, we can count on it to “work” in the sense Davidson requires.

³⁹⁰ Bilgrami 1992a, 80.

rather than ascriptional individualism.³⁹¹ If their defending ascriptional individualism is just an aberration,³⁹² why do they so often do so?

One reason would be that the two types of individualism are occasionally not distinguished clearly enough and that, in defending one, they feel that they must defend the other. This is especially clear with Davidson. While Davidson is certainly right in thinking that we can give up the intention to follow community norms in particular cases,³⁹³ one suspects that he occasionally runs together such cases (where we deliberately flout the societal norms) with those where we are merely ignorant of them. In the former case we lack the intention to be interpreted as others are, while in the latter case such an intention³⁹⁴ is probably present, and this difference is enough to make it the case that, even if the first type of speaker should be interpreted individualistically, the latter type of speaker usually should not.³⁹⁵ While there are occasions where we are willing to be interpreted idiosyncratically if our usage departs from the social norm, the intention to be interpreted as others are usually is the default case. Davidson occasionally suggests that just the opposite is true: that a willingness to be interpreted idiosyncratically should be understood as the default. It should, however, be clear that, in most contexts the default assumption should be that we will be interpreted according to standard usage, and so unless we are willing to flag a willingness to depart from such usage, we should intend to be interpreted in such a fashion. This phenomenon may be even more pronounced in certain more ‘disciplined’

³⁹¹ Indeed, in correspondence, Bilgrami explicitly endorses methodological over ascriptional individualism: “Bert has his own criteria beliefs about arthritis in each locality he uses it. Sometimes he may have in mind to say it about something that causes hurting in his thigh in the mornings, sometimes he may use it to be talking about whatever the doctor talks about when he uses it etc.” (Personal correspondence, July 1995)

³⁹² Though they both come out of a tradition that seems pretty clearly committed to ascriptional individualism. Quine’s original discussion of “radical translation” in *Word and Object* very explicitly ruled out interactions with “native kibitzers” whether the informant was disposed to rely on them or not.

³⁹³ These cases would presumably be limited to those words that refer to items for which we feel confident that we could set up adequate criteria without anyone’s help (I may be able to do this with “chair” but not with, say, “electron,” or, for that matter, “elm”).

³⁹⁴ Or at least the assumption/presupposition that we will be so interpreted.

³⁹⁵ Excepting those cases discussed in section 1.211 of this chapter (slips, malapropisms, and radical misunderstanding).

fields where there is a concerted effort to see to it that there is a clearly defined accepted usage (technical vocabulary) and participants in the field are expected to adhere to the accepted usage when using such vocabulary. We generally intend to be interpreted as our fellows are, and while such an intention is defeasible (if, for instance, our usage strays too far from the public norm),³⁹⁶ minor differences between ours and the public use will not be enough to make us give it up.

Davidson's concern about what is essential to language *as such* need not be relevant to what is, in fact, true for *our* language. As a result, the fact that language may turn out not to be *essentially* social need not affect the claim that contents *actually* ascribed to us should often be individuated non-individualistically. Davidson's interest in what is conceptually necessary for communication³⁹⁷ often leads him to overlook the fact that what is at issue with cases like Bert's use of "arthritis" is not what language *must* be like for *anything*, or even what it *must* be like for *us*, but rather how it is, *in fact*, for *Bert*.

Furthermore, both Davidson and Bilgrami interpret Burge as defending a position that is incompatible with even *methodological* individualism,³⁹⁸ and because of this, they naturally feel that Burge's arguments for his position need to be criticized. However, in arguing against Burge, both tend to contest the accuracy of Burge's claim about how Bert's utterance to his doctor should be interpreted, suggesting instead that Bert should here be understood as meaning something different than the doctor does by "arthritis." (In much the same way, they argue that Bert's deference should be explained not by his understanding himself as misapplying his term, but rather by his understanding himself as meaning something different by "arthritis" than his fellows and changing what he means to

³⁹⁶ What counts as 'too far' may be context-relative and determined in part by the judgments of the speaker.

³⁹⁷ Davidson is interested in "the *essential* nature of the skills" that are "*basic* to linguistic communication" and not "usual, though *contingent*" features of such communication. (Davidson 1982, 278, 280 [italics mine].)

³⁹⁸ Whether they are correct in doing so will not be discussed here. (Though Burge 1989 makes their reading of Burge considerably less plausible than Burge 1979 would have on its own.)

allow him to communicate more effectively with them.)³⁹⁹ However, arguing against Burge in this way implicitly commits them to ascriptional individualism, since Bert's interaction with his doctor is fairly close to a best-case scenario for allowing methodological individualism to justify ascriptionally non-individualistic attributions. Because of this, if we must interpret *Bert* individualistically *here* it seems quite plausible that we will have to interpret *everyone* individualistically *everywhere*. As a result, their strategy for defending methodological individualism implicitly commits them to ascriptional individualism.

If one wants to defend methodological individualism, there is no need to quibble with Burge about whether or not Bert has made a false claim about arthritis. The methodological individualist can freely admit that Burge is right about Bert's utterance to his doctor. All they need to do is argue that (1) such ascriptions are compatible with methodological individualism since they can be explained by Bert's reliance on the doctor's use in that context, and that (2) in contexts where no such reliance is present, there is no reason to ascribe contents in a social-externalist fashion. A defender of methodological individualism need not challenge the *correctness* of Burge's characterization of Bert's utterance, he need only challenge its *generalizability*.

As a result, then, methodological individualism, which clearly is an essential part of the Davidsonian program, can accommodate the types of ascriptionally non-individualistic intuitions that Burge appeals to. Consequently, Davidson and Bilgrami need not understand Burge's work as the kind of threat they frequently treat it as being.

³⁹⁹ See sect 1.222 of this chapter for a more extended discussion of this. Methodological individualism leads to ascriptional individualism if we deny that the speaker could have (or attach much weight to) the sorts of beliefs that would build a reliance on others into his concepts (such as the belief that doctors know much more about arthritis than he does), and such an account of deference would go a long way towards justifying such a denial.

Appendix E: Ascriptions *de dicto* and *de re*.

The accounts of self-knowledge and behavioral explanation presented in this chapter rely heavily on the *de dicto/de re* distinction, and the existence of such a distinction should be fairly uncontroversial. Both those who see the distinction as corresponding to two types of belief,⁴⁰⁰ and those who see it instead as corresponding to two types of ascription,⁴⁰¹ share the assumption that belief predicates are ambiguous.⁴⁰² However, there are many ways of drawing the *de dicto/de re* distinction, and the purpose of this appendix is to clarify the version of the distinction that is presupposed in this chapter.⁴⁰³

The reasons for thinking that belief predicates are ambiguous should be clear when we consider an ascription such as:

(I) Oedipus believes that his mother is the queen of Thebes.

There seems to be a sense in which the ascription is clearly true, and a sense in which is clearly false. It can seem true because the woman who was, in fact, Oedipus' mother was believed by him to be the queen of Thebes. It can seem false because Oedipus himself would have vigorously and sincerely dissented from the suggestion "your mother is the queen of Thebes." Belief ascriptions thus seem to admit of two interpretations,⁴⁰⁴ and this would suggest belief predicates are ambiguous.

However, while this chapter relies on the *de dicto/de re* distinction at a number of crucial points, it will be argued below that the impression of ambiguity that one can get by focusing on cases such as (I) is just another example of how hard cases make bad law.

⁴⁰⁰ See, for instance, Kaplan 1968, Burge 1977.

⁴⁰¹ See, for instance, Searle 1983 (esp. pp. 196, 217), Dennett 1982, Brandom 1994 (esp. ch. 8).

⁴⁰² As Stich recently put it "just about everyone who has written on [the *de dicto/de re* distinction] in the last twenty-five years has shared Quine's view that belief predicates are systematically ambiguous." (Stich 1986, 119.)

⁴⁰³ It should be noted that the characterization of the distinction does not involve a focus on the types of 'epistemically strong' *de re* beliefs that are associated with perceptual judgments and demonstrative utterances.

⁴⁰⁴ Once again, leaving it open whether these two interpretations should be taken to correspond to two types of belief or two types of ascriptions.

When we consider our ascriptional practices as a whole, the suggestion that belief predicates are ambiguous has little plausibility.

Why this is so can be seen when we look at Brandom's recent and detailed characterization of the *de dicto/de re* distinction as one between two types of ascription. Brandom characterizes these two types of ascriptions and their relation to the *de dicto/de re* distinction as one between "two different styles in which the content of the commitment ascribed can be specified."⁴⁰⁵ Consider, for instance, the following two ascriptions of belief to Oedipus:

(I) Oedipus believes that his mother is the queen of Thebes.

(II) Oedipus believes that his mother is the queen of Corinth.

Both (I) and (II) are ascriptions that can truly be made, the first because Oedipus believes that Jocasta is queen of Thebes, and the second because he believes that Merope is queen of Corinth. Of course, while we would endorse the characterization of Jocasta as "Oedipus's mother," Oedipus would not, and while Oedipus would endorse the characterization of Merope as "Oedipus's mother," we would not. According to Brandom, this "social difference of deontic attitude turns out to be what determines whether the ascription is *de dicto* or *de re*."⁴⁰⁶ Ascriptions such as (I) in which we characterize Jocasta in terms that we, but not Oedipus, would accept, are *de re*, and ascriptions such as (II), in which we characterize Merope in terms that Oedipus (though in this case not we), would accept, are *de dicto*. As Brandom puts it:

The substitutional commitments that govern the expressions used to specify the content of the commitment ascribed can . . . be either attributed to the one to whom the doxastic commitment is ascribed, or undertaken by the one ascribing it . . . [Where the ascriber employs] only commitments the ascriptional target would acknowledge, the content specification is *de dicto*. Where the ascriber has employed substitutional commitments he himself, but perhaps not the target, endorses, the content specification is *de re*.⁴⁰⁷

The *de dicto/de re* distinction is thus explained in terms of whether or not the content is ascribed in terms that either the ascriber or ascribee would accept. Such an account makes

⁴⁰⁵ Brandom 1994, 503.

⁴⁰⁶ Brandom 1994, 503.

⁴⁰⁷ Brandom 1994, 503, 507.

it trivially true that *de dicto* belief ascriptions will capture the ascriber's point of view, while letting the *de re* ascriptions capture what the ascriber takes this point of view to be a point of view on. Because of this, the fact that what we refer to may be determined by factors of which we are not aware does not threaten our knowledge of our beliefs as ascribed *de dicto*, or allow us to ascribe inconsistent beliefs to an agent in *de dicto* ascriptions such as "Oedipus believes that his mother is not his mother."

Still, while it seems promising to distinguish the *de dicto* from the *de re* in ascriptional terms, this *distinction* cannot be understood as a *dichotomy*. The distinction does not reflect an *ambiguity* in our belief predicates. To see why, we should first notice that, on such an account, many, if not most, ascriptions qualify equally as *de dicto* and *de re*. For instance, suppose I make the ascription:

(III) Oedipus believes that his daughter is good to him.

The characterization of Antigone as Oedipus's daughter is both one that I endorse and one that Oedipus would acknowledge. Ascriptions such as (III) would seem to involve both undertaking and attributing the commitment to the substitution in question, and thus might seem to count equally as *de re* and *de dicto*,⁴⁰⁸ which is why (III), like most attitude ascriptions, is not comfortably classified as just one or the other.⁴⁰⁹

How the ascriptional account can explain why most actual ascriptions seem neither purely *de dicto* nor purely *de re* can be shown by making clear how the two types of substitutional commitments (those accepted by the ascriber, those accepted by the ascribee) divide the possible substitution instances of a term into four classes.

1. Those acceptable to both the ascriber and the ascribee.
2. Those acceptable to the ascriber but not the ascribee.
3. Those acceptable to the ascribee but not the ascriber.
4. Those acceptable to neither the ascriber nor the ascribee.

⁴⁰⁸ Though, according to the passage from Brandom quoted above, they would count only as *de dicto*.

⁴⁰⁹ This is a distinct advantage the ascriptional account may have over those that treat the *de re/de dicto* distinction as one between two types of belief, since such accounts suggest that the belief in question must be one or the other (unless, of course, one is willing to say that there are two beliefs involved). (Ascriptional accounts also seem to deal better with iterated ascriptions. (See Brandom 1994, 608-613).)

Ascriptions of the first sort seem to be neither purely *de re* nor purely *de dicto*.⁴¹⁰ Ascriptions of the second sort could be characterized as ‘pure’ or ‘paradigmatic’ *de re*, while those of the third sort could be characterized as ‘pure’ or ‘paradigmatic’ *de dicto*. Belief ascriptions involving substitutional commitments of the fourth sort seem to be neither *de re* nor *de dicto*, but while they may not be very popular, there can be such ascriptions. In particular, there might be a substitution accepted by the ascriber’s audience (or some other salient group) but which the ascriber wants to make clear that neither he nor the ascribee accepts. Such a case might be an ascription such as “Chomsky believes of those ^sfreedom fighters^s that they are ruthless terrorists.”⁴¹¹

While it seems clear that most of our ordinary ascriptions are of type (1), cases of misidentification involve ascriptions of types (2) and (3).⁴¹² For instance, if Nicole shoots a cow thinking that it was a deer,⁴¹³ we may characterize her beliefs with sets of ascriptions such as:

Nicole believes of a cow that it is a deer. She thinks that, if she bags it, the ^sdeer^s will provide her with venison for the winter. She hopes that the ^sdeer^s won’t suffer much, and doesn’t want to hit the ^sdeer^s in the head, because she wants to mount the head on her wall.

Since the belief ascriptions most discussed in the philosophical literature (Superman cases, Pierre’s puzzle etc.) often gravitate to types (2) and (3), it can be easy to lose sight of the fact that most of our belief ascriptions involve characterizations of the topic that are accepted by *both* the ascriber and the ascribee.⁴¹⁴ Because of this, we should not claim, as

⁴¹⁰ Though, once again, Brandom’s official line seems to count both cases 1 and 3 as *de dicto*. However, on his regimentation, a *de dicto* ascription only commits one to being a case of type 3.

⁴¹¹ ^sScare quotes^s are here used to mark the fact that the ascriber does not accept the substitution in question, and they will be used more here than the “so called” locution for terminological convenience.

⁴¹² Of course, not all ascriptions of type (2) and (3) (especially *de re*) need involve mistakes.

⁴¹³ See Brandom 1994, 522ff.

⁴¹⁴ Proper names help ensure that most ascriptions will be of type (1), since they provide a type of common currency that makes such ascriptions easier. I can ascribe to my brother the belief that Cervantes wrote *Don Quixote*, and, while this ascription may turn out to be false, it is highly unlikely that it will be because, while he believes of Cervantes that he wrote *Don Quixote*, he does not believe it of him as “Cervantes.” There are, of course, people we have beliefs about without knowing their names (X believes of Orcutt that he is a spy), and cases where people have more than one name, but, in most cases, the use of a proper name will insure that the ascription will be of type (1).

Brandom does, that belief ascriptions are ambiguous.⁴¹⁵ If all ascriptions were either of types (2) or (3), such a claim would be uncontroversial: there would be just two types of ascriptions.⁴¹⁶ However, ascriptions of type (1) are not as easily viewed as being in any way ambiguous; they can express *both* a doxastic commitment and a doxastic attitude.

Since most ascriptions of belief are of type (1), one might argue that types (2) and (3) are (in spite of being ‘pure’ *de re* and *de dicto*) actually degenerate forms of belief ascription. In our unregimented language, cases of type (1) are the default cases. Unless we specify otherwise,⁴¹⁷ it is assumed that both the ascriber and the ascribee accept the substitutions in question. This is why (if one is unfamiliar with the story) statements like (I) and (II) can be misleading: our default assumption will be that both the ascriber and the ascribee are willing to endorse the characterizations in question.⁴¹⁸ Such statements would not be so misleading if there were a mere ambiguity involved, since there would be no such default assumption.⁴¹⁹ As a result, everyday belief ascriptions should not be seen as ambiguous between *de dicto* and *de re* ascriptions, but rather, pure *de dicto* and *de re* ascriptions should be understood as special cases in which we explicitly disavow some of the commitments associated with the everyday ascriptions. In particular, with *de dicto* ascriptions, we disavow any commitment to the substitution in question, while with the *de re* ascriptions, we disavow any commitment to the ascribee’s accepting the substitutions in

⁴¹⁵ “‘Belief’ is ambiguous in scorekeeping terms, referring sometimes to deontic status and sometimes to deontic attitude -- sometimes to doxastic commitment and sometimes to acknowledgment of such a commitment.” (Brandom 1994, 260.)

⁴¹⁶ Brandom’s “regimentation” of English is suited for such a two way ambiguity. Cases of type (2) appear in Brandom’s regimented language as “S believes *of t* that F(*it*),” while those of type (3) appear as “S believes that F(*t*),” but there is nothing in the regimentation to correspond to types (1) or (4).

⁴¹⁷ Or the context makes it obvious, as in “Grant believed that the site of the 1996 summer Olympics had to be captured.”

⁴¹⁸ We can thus bring the semantics of such terms comparatively close to our actual intuitions about their truth conditions. Both Brandom’s (1994) attempt to make the *de dicto* ascriptions the default and, say, Salmon’s (1986) attempt to characterize all such ascriptions as *de re* must rely to a much greater extent on ‘pragmatic’ considerations to explain our (apparently mistaken) intuitions about the truth conditions of these ascriptions.

⁴¹⁹ I.e., for paradigmatically ambiguous words such as “bank,” even if one of the two meanings has a certain ‘default’ status in various contexts (“I keep my money in the bank” etc.) it isn’t the case that the default status ever picks out *both* of the putative meanings. We don’t naturally assume that the speaker has in mind a financial institution on the side of a river.

question. These disavowals can be explicitly flagged with the “believes of” locution in the case of *de re* ascriptions and the “so-called” locution or ‘scare-quotes’ in the case of *de dicto*. The simple “X believes that P” locution is reserved for ascriptions of type (1), where the characterization is acceptable both to the ascriber and the ascriber.

In response to this, someone who wanted to keep the independence of our *de dicto* and *de re* ascriptions might argue that, rather than characterizing ascriptions of types (2) and (3) as explicit disavowals from the commitments involved in ascriptions of type (1), we can take ascriptions of type (1) simply as a *combination* of the commitments involved in ascriptions of types (2) and (3).⁴²⁰ The ‘pure’ *de dicto* and *de re* ascriptions would be primitive, and ‘mixed’ ascriptions of type (1) would be taken to be what you get when you put the two pure ascriptions together. However, there are at least two reasons to think that we should treat cases of type (1) as more primitive than those of types (2) or (3). The first is that our everyday ascriptions treat cases of types (2) and (3) as “marked”⁴²¹ cases while treating cases of type (1) as the default case. The second is that treating cases of type (2) and (3) as primitive makes it difficult to account for the possibility of cases of type (4). If one starts with just ascription-types geared towards ascribing the two types of commitment, it is hard to see how one is going to be able to use them to get a case where neither type of commitment is ascribed. On the other hand, if one takes cases of type (1) as primitive and cases of types (2) and (3) as produced by mechanisms that strip cases of type (1) of their usual commitments, then it should seem quite simple how one could get cases of type (4) by applying both mechanisms. Indeed, it would actually be a problem for this type of account if cases of type (4) were not possible.

⁴²⁰ This is the line Brandom suggests when he talks about how the cases in which ordinary that-clauses both characterize the speaker’s conception and pick out the relevant object can be regimented as, say “Alfred believes *of* the man in the corner *as* the man in the corner that he is a spy.” (Brandom 1994, 546.)

⁴²¹ “Believes-of” locutions and scare-quotes (corresponding to the colloquial “so-called”) respectively.

To return to the four types of substitutional commitments that can be expressed in a belief ascription, the suggested notation for them would be:

1. Those acceptable to both the ascriber and the ascribee. John believes that Peter is bald.
2. Those acceptable to the ascriber but not the ascribee. John believes of Peter that he is bald.
3. Those acceptable to the ascribee but not the ascriber. John believes that ^SPeter^S is bald.
4. Those acceptable to neither the ascriber nor the ascribee. John believes of ^SPeter^S that he is bald.

Since this regimentation allows one clearly to represent all four cases (while attaching the normal form to ascriptions of type (1) and registering the degenerate character of types (2) & (3)), it is more recognizably English than Brandom's (which attaches the normal form to the *de dicto*,⁴²² and which only allows for two types of ascriptions).⁴²³

Such an account, in which the 'pure' ascription forms, while acceptable, are in some sense degenerate, may suggest that belief predicates, rather than being strictly ambiguous are governed by something like a preference rule system.⁴²⁴ Examples of preference rule systems can be seen in the application of verbs such as *climb*. Even though climbing typically involves both (a) upward motion and (b) effortful grasping motions ("Bill climbed (up) the mountain"), neither the first nor the second condition need be satisfied ("Bill climbed down the mountain"; "The snake climbed up the tree"), though at least one of the two must be (* "The snake climbed down the tree"). This behavior is not best explained by suggesting that verbs like *climb* are ambiguous. The standard presupposition when we come across the verb is that *both* of the conditions are satisfied. Belief ascriptions could

⁴²² Brandom's mode of regimentation suggests that most actual belief ascriptions are *de dicto*, since it reserves the 'unmarked' "believes-that" locution for *de dicto* ascriptions, while the account here suggests that they are neither purely *de dicto* nor *de re*. Once again, Brandom does characterize cases of type (1) as *de dicto*; since *de dicto* attributions only involve attributing a substitutional commitment, his regimentation suggests that in everyday cases we undertake no commitment to the substitution in question.

⁴²³ *De re* ascriptions appear in Brandom's regimented language as "S believes *of t* that F(*it*)," while *de dicto* ascriptions appear as "S believes that F(*t*)." As a result, if the 'regimented' forms correspond only to the 'pure' *de dicto* and *de re* ascriptions, then there is nothing in the regimentation that allows one to specify a case as being of class (1) or (4), while if they pick out the 'mixed' ascriptions (i.e. both (1) and (3) count as *de dicto* and both (1) and (2) count as *de re*) not only is there still no way to pick out cases of type (4) at all, but there is no way to specify that one has cases such any of the other three types (i.e., (1) as opposed to (1) or (2)) in mind.

⁴²⁴ For a discussion of such "preference rule systems," see Jackendoff 1990, 35-6. (See also Lakoff 1987.)

be viewed as working in much the same way. The healthiest and ‘prototypical’ belief ascriptions involve substitution instances accepted by both parties, though attributions can still be made with substitution instances accepted by just one of the parties.⁴²⁵

⁴²⁵ This would suggest that cases of type (4) would not be acceptable unless explicitly marked by the grammar or context. So “Peter believes that the inventor of bifocals was the first postmaster” would be maximally acceptable if both Peter and the ascriber thought Franklin invented bifocals, moderately acceptable if just one of us did, and not acceptable at all (or at least no more acceptable than “the snake climbed down the tree”) if neither of us did.

Appendix F: Internalism, Individualism and the Background

Given that this chapter attempts to account for and justify ascriptionally non-individualistic practices within a methodologically individualistic framework, it is worth examining Searle's similar attempt to incorporate the truth of 'externalist' ascriptions within his 'internalist' framework. By discussing some of the virtues and faults of Searle's account, I will hope I will make the position defended in this chapter clearer.

Rather than challenging the intuitions Kripke and Putnam appeal to, Searle proposes to accept such intuitions and argue that they can be integrated into an 'internalist' framework.⁴²⁶ Indeed, rather than treating them as counterexamples to his position, Searle takes Kripke's, Putnam's and Donnellan's examples to illustrate how internalistic contents should be specified.⁴²⁷ Searle characterizes his position as "Holistic Internalism," which he defines as follows:

What is inside the head is entirely sufficient to determine the identity of each of our intentional states. The various "causal" and "contextual" conditions referred to by the externalists are entirely represented by the mind. However, the conditions of satisfaction of each intentional state are only fixed relative to the Network and the Background which is not part of each intentional state though, of course, the Network and the Background are nonetheless parts of the mind. (Searle 1991, 237.)

Searle claims that critics of internalism have missed this alternative because they "look only at what the agent might *say* and not to the *total Intentional content* he has in his head, and also because they neglect the role of the Network and the Background."⁴²⁸ Searle's suggestion that various 'background' capacities would allow one to justify ascriptionally non-individualistic content ascriptions is potentially very fruitful. Unfortunately, it is not one that Searle fully exploits.

To see why, we need to consider some of what Searle says about the "Background." The Background is crucial for Searle because "in general, representations only function, they only have the conditions of satisfaction that they do, against this nonrepresentational

⁴²⁶ Searle 1983, 203.

⁴²⁷ Searle 1983, 204, 206, 244-6.

⁴²⁸ Searle 1983, 250.

background.”⁴²⁹ This Background is “a bedrock of mental capacities that do not themselves consist in Intentional states (representations)”⁴³⁰ but which includes “various skills, abilities, preintentional assumptions and presuppositions, stances, and nonrepresentational attitudes.”⁴³¹ Furthermore, Searle is quite explicit that not only should elements of the Background not be viewed as intentional states themselves, but also they should not be seen as *part* of the content of our intentional states.⁴³² As a result, even if we believe that facts about what we mean must ultimately be explained in terms of facts about what is going on inside our heads, if some of these facts are part of the Background, then they need not be built into the Intentional content. They affect the Intentional content by helping to determine the conditions of satisfaction, but they do not enter into a specification of the Intentional content itself.⁴³³

⁴²⁹ Searle 1983, 21.

⁴³⁰ Searle 1983, 143.

⁴³¹ Searle 1983, 151. The following quotations should give a rough feel for how Searle understands the Background:

If we actually tried to carry out the task [of specifying each of the intentional states in the Network] we would soon find ourselves formulating a set of propositions which would look fishy if we added them to our list of beliefs in the Network; “fishy” because they are in a sense too fundamental to quality as *beliefs*, even as unconscious beliefs. Consider the following propositions: elections are held at or near the surface of the earth; the things people walk on are generally solid; people only vote when awake; objects offer resistance to touch and pressure. (Searle 1983, 142.)

A man certainly *could* have the belief that tables offer resistance to touch, but ... that isn’t the correct way to describe the stance that I, for example, now take towards this table and other solid objects. For me, the hardness of tables manifests itself in the fact that I know how to sit at a table, I can write on a table, I put stacks of books on tables ... as I do each of these things I do not in addition think unconsciously to myself, “it offers resistance to touch.” (Searle 1983, 142.)

⁴³² Searle 1983, 19, Searle 1991, 230, 232-3.

⁴³³ As a result, one need not accept Searle’s claim that “Even if I am a brain in a vat ... nonetheless, I do have the Intentional content that I have, and thus I necessarily have the same Background that I would have if I were not a brain in a vat and had that particular Intentional content.” (Searle 1983, 154.) One reaches such a conclusion only if one builds the Background into the Intentional contents themselves, rather than simply allowing that the Background helps determine what the Intentional contents are. I would have different Intentional contents if I were a brain in a vat, because the Background which the brain and I share produces different Intentional contents in the different contexts. (Incidentally, Searle’s claim that I would have the same Intentional contents even if I were a brain in a vat points to a serious weakness in his discussion of the twin cases. Namely, he accounts for the fact that I and my twin-earth counterpart have different Intentional contents by the fact that we are different people and thus have different thought/experience tokenings. His account is thus not able to distinguish Intentional contents if we think of the *same* experience as arising in different counterfactual situations (as his discussion of the brain in the vat makes abundantly clear). As a result, just as Searle says that I would have the very same Intentional contents were I a brain in a vat, he must also say that I would have the very same

Unfortunately, Searle himself often fails to “distinguish those features of content which function to determine conditions of satisfaction from features outside the content which also serve to fix conditions of satisfaction.”⁴³⁴ In particular, Searle’s specification of the conditions of satisfaction of our intentional states in causally self-referential terms seems to build aspects of the Background into the Intentional content itself.⁴³⁵ As a result, in spite of his insistence that the Background is “non-representational,”⁴³⁶ he goes on to claim that “the various ‘causal’ and ‘contextual’ conditions referred to by externalists are entirely *represented* by the mind.”⁴³⁷ Searle thus seems to miss his own point that while “it is always possible to take an element of the Background and treat it as a representation ... from the fact that it is possible to treat an element of the Background as a representation, it does not follow that, when it is functioning, it is functioning as a representation.”⁴³⁸

For instance, it is plausibly part of the Background that we understand ourselves as perceiving objects in an independent world that causally impinges upon our senses.⁴³⁹ This understanding of our place in, and relation to, the world we perceive is manifested in our everyday behavior and can be understood as part of the preintentional Background. Because of this, it is quite plausible to think that the “causally self-referential” character of visual perception is a function of the Background, but Searle builds it directly into the perceptual content itself, claiming that the intentional content of, say, my visual experience of seeing a yellow station wagon before me should be spelt out as:

Intentional contents if I were brought up on Twin Earth instead of Earth, in spite of the fact that the people etc. I interact with are completely different.)

⁴³⁴ Searle 1991, 232.

⁴³⁵ Searle admits that he occasionally makes such an assimilation, but he is clear that such assimilations are only a type of “shorthand”: “When arguing in favor of ‘internalism’ . . . I sometimes use the expression ‘Intentional content’ as a shorthand for Intentional content including Background and Network conditions on its conditions of application, all of which are internal to the mind.” (Searle 1991, 233.)

⁴³⁶ Searle 1983, 21.

⁴³⁷ Searle 1983, 237, italics mine.

⁴³⁸ Searle 1983, 158.

⁴³⁹ This might be related to the commitment to ‘realism’ that Searle takes to be part of the Background. (Searle 1983, 159.)

I have a visual experience (that there is a yellow station wagon there and that there is a yellow station wagon there is causing this visual experience).⁴⁴⁰

Searle's specification of the Intentional content thus ignores a possibility that his account clearly allows for: the content of the visual experience involves the yellow station wagon itself. Our perception has this content only thanks to elements in the Background, but these background elements are not part of the content. It need not follow from the fact that one must be able ultimately to give a first-personal account of our semantic concepts that our Intentional contents themselves should be specified in such an explicitly first personal manner.

While Searle's actual account of Intentional contents seems unsatisfactory, his general approach of stressing the importance of the background implicit in our practices can be extended in a number of productive ways. For instance, the Background can include not only a general picture of how we are causally related to the objects of perception, but also our 'implicit'⁴⁴¹ understanding of the 'informational network' we are part of. That is to say, while we may not explicitly think that, say, our memories refer to people who are causally responsible for those very memories, such an 'assumption' is implicit in our practices, and forms part of the Background against which our thoughts get the contents they have. If we remember something, we assume we were there, and if we discover that the person we take a memory of ours to be a memory of could not have been causally responsible for the memory, then we assume that we must have been thinking of someone else. As a result, a subject can be said to 'know which' object he is thinking of even if he cannot explicitly *state* the object's role in the informational system.

With these background capacities in place, one can satisfy Russell's Principle ("a subject cannot make a judgment about something unless he knows which object his

⁴⁴⁰ Searle 1983, 48. Such an account of the content of our perceptual experiences is not meant to correspond to anything running through our mind when we have such experiences, rather, it represents "a verbal specification of what the Intentional content requires if it is to be satisfied." (Searle 1983, 49.)

⁴⁴¹ Such talk of our understanding of something as 'implicit' should in no way be understood as suggesting that what is understood as 'represented' in any way. Rather, it only suggests that we will act in a way that will manifest a commitment to the truth of what is implicitly understood.

judgment is about’’) ⁴⁴² in cases where Evans might not have originally thought it possible. For instance, Evans gives the following as an example of someone who could not satisfy Russell’s principle, and thus not have the thought about a particular shiny ball.

Our subject briefly sees one ball rotating by itself on one day, and the other on a later day. And let us further suppose that the subject retains no memory of the first episode, because of a localized amnesia produced by a blow to the head. Suppose, finally, that many years later our subject reminisces about ‘that shiny ball’ he saw many years earlier. If asked which ball he is thinking about, our subject cannot produce any facts which would discriminate between the two. ⁴⁴³

Evans admits, of course, that the subject could specify the ball in question if he reflected on the role of memory in the informational system, but the resulting thought would be of an entirely different kind.

It is true that *if S reflects on the process of information-transmission*, he may, so to speak, aim a thought (a particular-thought) at ‘that ball.’ But the object of the thought will be individuated, in this case, *by reference to its role in the operations of the information system*. This is a different kind of thought from the kind of particular-thought which it is the aim of this chapter to discuss. ⁴⁴⁴

Furthermore, he claims that “it is surely correct to assume that thinkers (among whom we have to include very young children) will not in general resort to this [causal] way of identifying the object of their thought.” ⁴⁴⁵ However, while the typical subject will not *appeal* to the object’s role in the informational system, there is still a sense in which the subject *knows* that the object must play such a role, and can thus be said to know which object he is thinking of. Particular-thoughts may be of a different kind from those that make *explicit* reference to the object’s role in the informational system, but there is no reason not to think that such thoughts may still rely on an *implicit* awareness of such a role to get their content. ⁴⁴⁶ When background capacities are brought into play, Russell’s

⁴⁴² Evans 1982, 89.

⁴⁴³ Evans 1982, 90.

⁴⁴⁴ Evans 1982, 128. (The particular example has been changed from ‘that native’ to ‘that ball’ for continuity.) See also, “It is an assumption of the case that the subject would not have any notion of appealing to these causal considerations; otherwise the case would cease to be a counterexample to Russell’s Principle” (Evans 1982, 117).

⁴⁴⁵ Evans 1982, 117. See also: “I assume, as I think that I am entitled to, that he would not think of distinguishing the ball he is thinking of as the one from which his current memory derives.” (Evans 1982, 117.)

⁴⁴⁶ Evans is surely right to think that thoughts that bring these background capacities *explicitly* into play are of a different type from those that do not, and if this is the case, then one has another reason for thinking that, *pace* Searle, characterizations of the Intentional contents of our thoughts relying on such a background should not make explicit reference to their ‘self-referential’ character.

principle seems much less restrictive. We can allow that *merely* being causally connected to an object is not enough to have thoughts about it (which is the gist of the “photograph model” Evans sees as in opposition to Russell’s principle), while also insisting that our implicit understanding⁴⁴⁷ of the informational system helps constitute our ability to think of those items from which our memories derive.

The sorts of ‘intuitive’ considerations Kripke and Donnellan muster in support of their accounts can be understood as pointing to the fact that we must have something like a background capacities that can be understood in terms of the informational system. If such capacities were not in the Background, there is no reason that we would find it ‘obvious’ that we are not referring to some causally isolated hermit by “Thales,” no matter how many of our Thales-beliefs are true of him. Indeed, it is just such reactions that constitute our having a commitment to the information system. As a result, we could argue that everybody already ‘knew’ some version of Kripke’s account, and they just needed his help, through a *Meno*-like process, to ‘recollect’ it.

A similar set of Background capacities involving how we understand the language we share with others can be taken to be implicit in our practice. Not only do we have an implicit understanding of how perception and memory are causally related to the world around us, we have an implicit understanding of how language and linguistic communication function, an understanding that involves our speaking shared languages with shared meanings. Our implicit commitment to this model is manifested in not only our content attributions, but also our deference behavior.⁴⁴⁸ For example, when our use is corrected, our feeling of having made a certain sort of mistake (misusing a word, saying something false (as opposed to using a non-standard word and saying something

⁴⁴⁷ I.e., a set of practices that could be explicitly represented as such a system.

⁴⁴⁸ If socially infected content ascriptions are justified by the presence of such a background picture of language and such a background picture is manifested in, say, the type of deference behavior described above, then it is of considerable importance that we do defer in such a fashion. Explanations of deference in terms of, say, a pragmatic desire to communicate effectively, if true, would suggest that there is no such background picture of a shared language, in which case it would be extremely difficult to avoid some type of ascriptional individualism. These issues are discussed in fuller detail in 1.222.

misleading))) is explained by, and is partially constitutive of, our having such a picture of language in the background. Kripke's claim that "when the name 'is passed from link to link,' the receiver of the name must ... intend when he learns it to use it with the same reference as the man from whom he heard it"⁴⁴⁹ can be understood as a somewhat voluntaristic characterization of aspects of this Background understanding of language. With this Background in place, many 'non-individualistic' ascriptions are subject to methodologically individualistic (and in some sense implicitly metalinguistic) explanation without running into the type of objection that Burge lodges against straight metalinguistic accounts. Namely:

There appears to be a general presumption that a person is reasoning at the object level, other things being equal. The basis for this presumption is that metalinguistic reasoning requires a certain self-consciousness about one's words and social institutions. This sort of sophistication emerged rather late in human history. (Cf. any history of linguistics.) Semantical notions were a part of this sophistication. (Burge, 1979, 97)

The idea is *not* that the intentional content of my "arthritis" belief is spelled out in terms of whatever the community which is causally responsible for my use of "arthritis" means by the term (as it would be on the simple metalinguistic account). Rather the idea is that the intentional content simply involves arthritis, but that it does so in virtue of a non-intentional background that implicitly has a picture in place that has such a result as its consequence.

Searle, of course, does not see things this way and tries to put the Background directly into explicitly metalinguistic intentional contents when he claims that "Often, in fact, one does make what I called parasitic reference using a proper name: often the only identifying description one associates with a name '*N*' is simply the 'object called *N* in my community or by my interlocutors.'"⁴⁵⁰ In just the way his account of the content of visual experience did, these formulations neglect the point that the Background should be in the *background*. Such formulations, in addition to being implausibly metalinguistic, are susceptible to various types of counterexamples (e.g., the generally accepted usage has changed since I learnt the name). Much the same can be said of his attempt to deal with the Putnam cases

⁴⁴⁹ Kripke 1972, 96.

⁴⁵⁰ Searle 1983, 243. See also Searle 1983, 244, 250.

with claims such as “The indexical definitions given by Jones on earth of “water” can be analyzed as follows: “water” is defined indexically as whatever is identical in structure with the stuff causing *this* visual experience, whatever that structure is.”⁴⁵¹ This, once again, builds too much Background into the content and is also subject to obvious counterexamples.

⁴⁵¹ Searle 1983, 207.