



## The Methodology of Naturalistic Semantics

Michael Devitt

*The Journal of Philosophy*, Vol. 91, No. 10 (Oct., 1994), 545-572.

Stable URL:

<http://links.jstor.org/sici?sici=0022-362X%28199410%2991%3A10%3C545%3ATMONS%3E2.0.CO%3B2-6>

*The Journal of Philosophy* is currently published by Journal of Philosophy, Inc..

---

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/jphil.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

---

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).

## THE METHODOLOGY OF NATURALISTIC SEMANTICS\*

Semantics is a veritable Balkans of the intellectual world. Localists war with holists, truth conditionalists with verificationists, deflationists with substantivists, direct-reference theorists with Fregeans, one-factor theorists with two-factor theorists, and so on. An army of enthusiasts for narrow content have occupied the territory formerly held by the proponents of wide content. Finally, no settlement of these disputes seems to be in sight.

One sound stands out in these battles: the clash of semantic intuitions. Indeed, sometimes that is the only sound to be heard. Intuitions are almost always aired in "thought experiments."

This reliance on intuitions may be untroubling from some perspectives because it seems to exemplify the characteristic method of "armchair" philosophy. Yet it is surely troubling from the naturalistic perspective that I favor. According to naturalism, semantics is an empirical science like any other. Intuitions and thought experiments do not have this central role elsewhere in science. Why should they in semantics?

This question leads to the general ones that are the main concern of this paper: How should we get to the truth in semantics? How should we go about settling semantic disputes? What is the right methodology for semantics?

A naturalistic approach to these questions can only hope for modest answers. We cannot expect to make more progress with the methodology of semantics than has been made with scientific methodologies in general. And we know how limited that progress is. It has turned out to be very hard to say how we should get to the truth and settle disputes in science. My hope is only to bring semantic methodology close to other scientific methodologies.

In the next section, I shall say more about why the reliance on semantic intuitions is troubling. I move on to consider the question: What is the semantic task? To give a methodology for semantics, we

\* This paper has benefitted from discussions of earlier versions delivered at the Graduate Center of the City University of New York, Virginia Commonwealth University, Buffalo State College, the NEH Institute, "The Nature of Meaning," organized by Jerry Fodor and Ernie Lepore at Rutgers University, Franklin and Marshall, the Australian National University, Monash University, and Flinders University. I am also indebted to David Braun, Fiona Cowie, Peter Godfrey-Smith, Eric Lormand, Peter Ludlow, Greg Ray, Georges Rey, Michael Slote, and Corliss Swain for comments.

need an answer to this question and yet it is far from obvious what the answer is. I propose an answer and then a methodology. Finally, I make some predictions about the results of applying the methodology. Despite the modesty of the methodology, I think that it yields a speedy settlement of several, though certainly not all, semantic disputes.

#### I. THE TROUBLE WITH INTUITIONS

Clearly, there could be no objection to the standard appeal to intuitions in semantics if the semantic task were simply the systematization or explanation 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.

Reliance on semantic intuitions might also be appropriate if these intuitions were pieces of a priori knowledge. Many treat them as if they were. Yet from the naturalistic perspective, we should doubt that there is *any* a priori knowledge. Indeed, I think that we should doubt that any satisfactory explanation can be given of what a priori knowledge would be and of how it might be possible. And, even if there were some a priori knowledge, there clearly would not be enough of it to form an adequate evidential base for physics, biology, or economics. Why then would there be enough for semantics?

Many think that our semantic intuitions have some special authority because they reflect our linguistic (conceptual) competence. This competence is thought to give us "privileged access" to meanings. Merely understanding the words 'bachelor' and 'unmarried' (merely having the concepts "bachelor" and "unmarried") yields propositional knowledge of their meanings and of the relations these meanings have to each other. And this "tacit theory" can be brought before the conscious mind by "analysis." But why should we accept this vestige of Cartesianism? It often seems to be based on little more than the beguiling use of the vague term 'know' in the ordinary locution, 'knows the meaning'.<sup>1</sup> We should see our competence as simply an ability, a piece of knowledge-how not knowledge-that.

<sup>1</sup> The history of this semantic Cartesianism is strange. The most common version assumes that our competence consists in propositional knowledge of *truth conditions*. Herbert Heidelberger has shown how widespread this assumption is with references to David Wiggins, Peter Strawson, Donald Davidson, Gottlob Frege, Ludwig Wittgenstein, W. V. Quine, and Rudolf Carnap: "Understanding and Truth Conditions," *Midwest Studies in Philosophy, Volume V: Studies in Epistemology* (Minneapolis: Minnesota UP, 1980), pp. 401-10. He points out that it seems to be

Ordinary semantic intuitions should have the same epistemic status as ordinary intuitions anywhere. They are parts of an empirical, fallible, and certainly inadequate set of folk opinions or, more pretentiously, “folk theory,” the linguistic wisdom of the ages. What role then should they play in semantics? A major aim of this paper is to answer this question.

Many assume that the “armchair” thought experiments characteristic of philosophy show that it has a special, nonscientific way of getting knowledge. A by-product of my discussion of the role of intuitions is a naturalistic account of this philosophical method (section IX).

## II. THE TROUBLE WITH THE SEMANTIC TASK

If the semantic task is not to systematize or explain semantic intuitions, what is it? What is semantics trying to do?

There seems to be a simple answer: the “basic” task is to say what meanings *are*, to explain their *natures*. It is thus analogous to such tasks as saying what genes, atoms, acids, echidnas, or pains *are* but not, we should note, to such tasks as saying what genes, etc., *do*, stating the laws that advert to them. (The difference is illustrated by the difference between molecular and Mendelian genetics.) We start the semantic task, however, in rather worse shape than we do its ana-

---

regarded as “uncontroversial ... harmless ... perhaps unworthy of serious discussion” (p. 402). Gareth Evans says, “perhaps no one will deny it” — *The Varieties of Reference*, John McDowell, ed. (New York: Oxford, 1982), p. 106. Yet, as Heidelberger notes, it is not “obviously true.” He makes what, so far as I know, is the first attempt to argue for it. He does not claim success. I attempted to construct an argument based on hints in the literature but it turned into a travesty: *Realism and Truth* (Cambridge: Blackwell, 2nd ed. rev. 1991), pp. 270-2. For examples of the hints, see Colin McGinn, “The Structure of Content,” in *Thought and Object*, A. Woodfield, ed. (New York: Oxford, 1980), pp. 207-58; Crispin Wright, “Truth Conditions and Criteria,” *Proceedings of the Aristotelian Society*, Supp. Vol. L (1976): 217-45; and particularly Michael Dummett, *Truth and Other Enigmas* (Cambridge: Harvard, 1978), pp. 124-6. Apart from some brief and inconclusive remarks of Dummett in *The Interpretation of Frege’s Philosophy* (Cambridge: Harvard, 1981), pp. 310-1, that seems to be all that has been said for the assumption. Yet, despite this unjustified popularity, the assumption has received surprisingly little criticism: Gilbert Harman, “Language, Thought, and Communication,” in *Language, Mind, and Knowledge*, *Minnesota Studies in the Philosophy of Science*, Volume VII, K. Gunderson, ed. (Minneapolis: Minnesota UP, 1975), p. 286; my *Designation* (New York: Columbia, 1981), pp. 95-110; my “Realism and Semantics,” part II of a critical study of *Midwest Studies in Philosophy*, Volume XII, in *Noûs*, xvii (1983): 674-5 (on Heidelberger); my *Realism and Truth*, pp. 270-5; Scott Soames, “Semantics and Psychology,” in *The Philosophy of Linguistics*, J. Katz, ed. (New York: Oxford, 1985), pp. 204-26; G. Rey, “The Unavailability of What We Mean: A Reply to Quine, Fodor, and Lepore,” *Grazer Philosophische Studien*, XLVI (1993): 61-101. (Similarly, linguists often seem to identify syntactic competence with propositional knowledge of grammatical rules. Interestingly, this assumption has received more philosophical criticism.)

logues. With them, the subject matter of investigation is already identified relatively uncontroversially. This reflects the fact that we have clear and familiar theoretical or practical purposes *for which* we identify the subject matter. Semantics does not start out like that. It is far from clear what counts as a meaning that needs explaining. Indeed, the intractable nature of semantic disputes largely stems from differing opinions about what counts.<sup>2</sup> William Lycan<sup>3</sup> has brought out the problem wittily with his “double indexical theory of meaning”:

MEANING =<sub>def</sub> Whatever aspect of linguistic activity happens to interest *me now* (*ibid.*, p. 272).

We start semantics in the unusual position of having to specify a subject matter. We should not insist on great precision about this in advance of theory, but we do need some explication of our vague talk of “meanings.” And to avoid Lycan’s mockery, we must specify a subject matter *worthy* of investigation; we need an explication that is *not ad hoc*. Finally, the semantic task of interest to philosophers should be not only worthwhile but *fundamental*. To meet these needs, I propose to address three questions. First, what might plausibly be seen as our ordinary way of ascribing meanings? Second, to what do we ascribe them? Or, putting this another way, what are the phenomena that concern semantics? Third, what is our semantic interest in these phenomena? Or, putting this another way, what purposes are we serving in ascribing meanings?

### III. ASCRIPTIONS OF MEANING; SEMANTIC PHENOMENA

What might plausibly be seen as our ordinary way of ascribing meanings? Consider the following: “Ruth believes that Gorbachev has fallen” and “Adam said that Yeltsin has risen.” Such “propositional-attitude ascriptions” mostly use no semantic words but nevertheless seem partly to ascribe meanings. My *working assumption*, for this section and the next, is that they do indeed partly ascribe meanings: we specify meanings by the ‘that’ clauses in attitude ascriptions (also by their ‘to’ clauses that follow ‘want’, etc., and by clauses that follow ‘wonder whether’; briefly, by “*t*-clauses”).

<sup>2</sup> James Higginbotham, “Truth and Understanding,” *Iyyun, The Jerusalem Quarterly*, XL (1991), remarks: “The chief problem about semantics comes at the beginning. What is the theory of meaning a theory of?” (p. 271). Ned Block, “Advertisement for a Semantics for Psychology,” *Midwest Studies in Philosophy, Volume x: Studies in the Philosophy of Mind* (1986), notes: “Meaning is notoriously vague” (p. 615).

<sup>3</sup> *Logical Form in Natural Language* (Cambridge: MIT, 1984).

The meanings specified by *t*-clauses are complex. They are composed of simpler meanings like those specified by 'Gorbachev' and 'fallen'.

Two remarks about usage. I shall continue to talk of the "meanings" of thoughts where some would talk of their "contents." I shall adopt a usage according to which meanings are *properties*. This usage sits well with one popular view of the "logical form" of attitude ascriptions, which I favor, but less so with another.

It sits well with the view that the *t*-clause in an attitude ascription functions like an *indefinite* singular term. So just as 'Ken loves an echidna' asserts a relation between Ken and some object that has the property specified by 'echidna', 'Ruth believes that Gorbachev has fallen' asserts one between Ruth and some object that has the property specified by 'that Gorbachev has fallen'.<sup>4</sup> On this view, the latter property is naturally called "a meaning" and the object that would have it if the ascription were true would be concrete: a token thought. And if the ascription were to an utterance then the object that would make it true would also be concrete, a token utterance.

The usage sits less well with the view that the *t*-clause in an ascription functions like a *definite* singular term; for example, like a name. So just as 'Ken loves Gaelene' asserts a relation between Ken and the object designated by 'Gaelene', the belief ascription asserts one between Ruth and the object designated by 'that Gorbachev has fallen'. That latter object, "a proposition," is commonly called "a meaning."<sup>5</sup> In taking meanings to be properties, I depart from this usage. According to my usage, on this view of the logical form of ascriptions, the property of *having an attitude to* the proposition is called "a meaning." And the object that would have the property if the ascription were true would still be concrete.

My usage has two advantages. First, it avoids commitment to propositions. Second, whereas if the meaning of a token is an object then the token obviously has the property of being related to that object, it is not the case that if the meaning is a property then that property is obviously a relation to some object. My usage leaves the latter issue open.

<sup>4</sup> This view goes back at least to Wilfrid Sellars, *Science, Perception, and Reality* (New York: Humanities, 1963); see also, Davidson, *Inquiries into Truth and Interpretation* (New York: Oxford, 1984), pp. 93-108; Lycan, *Judgement and Justification* (New York: Cambridge, 1988), pp. 7-9.

<sup>5</sup> For two recent examples of this view, see M. Richard, *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them* (New York: Cambridge, 1990), p. 5; Stephen Schiffer, "Belief Ascription," this JOURNAL, LXXXIX, 10 (October 1992): 499-521.

We ascribe meanings to concrete thoughts and utterances. So, thoughts and utterances are the immediate phenomena of semantics. But thoughts, at least, are not so immediate to observation. What phenomena lead us to the view that an object has thoughts? Partly, the utterances of the object make us think this. Utterances are linguistic behavior. Clearly, the object's nonlinguistic behavior may also lead us to the view that it has thoughts.

#### IV. SEMANTIC PURPOSES

What is our semantic interest in these phenomena? What significant purposes — explanatory, practical, or whatever — are served by the ascription of meanings which, according to our working assumption, is part of the ascription of thoughts and utterances like those to Ruth and Adam?<sup>6</sup> Doubtless there are several such purposes but I shall focus on two: first, to explain and predict behavior, which I shall abbreviate “to explain behavior”; and, second, to use the thoughts and utterances of others as guides to a reality largely external to the subject. I shall consider these in turn. Our interest in thoughts is primary, our interest in utterances, secondary. So I shall start with thoughts.

(1) Consider this explanation of nonlinguistic behavior:

Why did Granny board the bus? She wants to buy a bottle of scotch. She believes that she has her pension check in her pocket. She believes that the bus goes to her favorite liquor store.

Such “intentional” explanations of behavior are familiar and central parts of ordinary life, of history, of economics, and of the social sciences in general. They all ascribe thoughts with meanings specified by *t*-clauses.

Consider this explanation of linguistic behavior next:

Why did Granny produce the sound /I need a drink/? She believes that she needs a drink. She wants to express her belief to her audience. /I need a drink/ expresses that belief.

Again, we ascribe thoughts to explain behavior. Such explanations are not so common because they are so obvious. They are implicit in our responses to communications.

(2) Ascribing beliefs serves another remarkably valuable purpose. If a person believes that the world is such and such, and is reliable, then we have good reason to believe that the world is such and such. Thus, attributing the property of meaning “it is raining” to Mark's

<sup>6</sup> I have made several attempts to answer this question in the past, most recently in *Realism and Truth*, ch. 6. These answers all have similar elements but I think that none of them has been right.

belief not only helps to explain his rain-avoidance behavior but also gives us evidence about the weather. We can even learn from someone who is a reliable guide to the way some area of the world is *not*.

Turn now to the ascription of *utterances*. Granny's behavior leads us to remark, "Granny says that she needs a drink." If the English-speaking Mark produces /It is raining/, we may respond, "Mark says that it is raining." If we think that Granny's and Mark's utterances are sincere expressions of beliefs, we shall ascribe *the same* meanings to their beliefs as we do to their utterances. Utterances are indicative of thoughts. Indeed, it is *because of* our interest in thoughts that we are interested in utterances. Thus, it is because we want to use Granny's thoughts to explain her behavior that we are interested in her utterance. And it is because we want to use Mark's thoughts to explain his behavior and inform us about the weather that we are interested in his utterance.<sup>7</sup>

Return to our use of others as guides to reality. We have a wide range of interests in learning about the world. The direct way to serve these interests is to examine the world. The indirect way is to use indicators. Sometimes these indicators are "natural" ones like tree rings. Sometimes they are artifacts like thermometers. Very often they are the beliefs of others. Some belief ascriptions serve our *theoretical* interest in explanation. Many, however, are like ascriptions of desires, hopes, etc., in serving interests that are not really theoretical at all. We have the most immediate *practical* interest in finding out quite humdrum facts about the world to satisfy our needs for food, shelter, a mate, and so on. So it helps to know what is on sale at the supermarket, where there is a hotel, who is available, and so on. Ascribing beliefs is a very good way of finding out about anything at all.

Humans gain most of their knowledge — what we learn at mother's knee, from teachers, and from books — by this process of attributing beliefs. The significance of this for the human species could hardly be exaggerated: we have a way, other than the slow and painful one through the genes, of passing on the discoveries of one generation to the next.

Or so it seems. But maybe we are all *wrong* in attempting to learn from others in this way. Maybe we are also wrong in attempting to

<sup>7</sup> The story is a bit more complicated because of the Gricean distinction between the speaker meaning of an utterance and its conventional meaning on the occasion. It seems that we ascribe the latter with the "says that" construction, yet we identify the former with the thought meaning in the normal situation. If we want to draw attention to a difference between the two meanings it seems that we use the construction "says that ... but ... means that ..."



explain behavior in the earlier way. Perhaps we are totally mistaken and should become behaviorists or eliminativists of some other sort. Or perhaps we should be less radically revisionist, ascribing only "narrow contents" or syntactic properties to serve our purposes. Still, my hypothesis is that, rightly or wrongly, we do ascribe meaningful thoughts and utterances to serve these purposes.

In identifying these two purposes — explaining behavior and guiding us to reality — as ones for which we ascribe meanings I do not mean to suggest that there are not others. One other, for example, arises from our interest in a desirer: if we wish her well we may want to satisfy her desires; if we wish her ill we may want to frustrate them. My claim is only that the purposes identified are significant and that we expect our ascriptions of meanings to serve them, *at least*.

Doubtless these two purposes can be served in many ways. Our interest in meanings, together with the working assumption, have led us to two particular ways of serving them: directly, by ascribing to a token that is thought a property of the sort specified by *t*-clauses; indirectly, by ascribing such a property to a token that is uttered. Let us say that in attempting to serve the purposes in these ways, our purposes are *semantic*. And let us say that a property plays a semantic role if and only if it is a property of the sort specified by *t*-clauses and, if it were the case that a token thought had the property, it would be in virtue of this fact that the token can explain the behavior of the thinker or be used as a guide to reality.<sup>8</sup>

#### V. SEMANTIC TASKS

Earlier, I stated the "basic" task as: to explain the natures of meanings. This statement needed an explication that was not *ad hoc*; we needed to specify a subject matter worthy of study. To meet these needs I adopted a working assumption: we specify meanings by the *t*-clauses of attitude ascriptions. Guided by this I have described purposes and roles that I have called semantic. Now, it is certainly worthwhile to study a property which plays a semantic role, which plays one or both of the roles just outlined in explaining behavior and informing us about reality. Furthermore, the study of such properties seems appropriately fundamental. So, I propose to add the following explication to the statement of the basic task: a property is a meaning if and only if it plays a semantic role. We can now drop our working assumption: what we specify by a *t*-clause will count as a

<sup>8</sup> Perhaps this should be broadened to allow tokens in, say, a visual module — hence not objects of *thought* — to have a semantic role, but I shall not attempt to do so.

meaning if and only if it does indeed have a semantic role. Any other property will count as a meaning if and only if it has such a role.

It might be objected that this explication is too liberal because not just *anything* that had a semantic role would be a meaning. It is usual to think that meanings are constituted by relational properties: “internal” ones involving inferential relations among tokens and/or “external” ones involving certain direct causal relations to the world. So it might be claimed that we should modify our explication so that only properties with this “appropriate” sort of constitution count as meanings.

Nothing in my argument hinges on my not adopting this modification, but I shall not adopt it. Note that my definition of “semantic role” slipped in the following constraint: a property has a semantic role only if it is “of the sort specified by t-clauses.” So the “liberal” explication already places this constraint on meanings. In effect, the modification makes this somewhat vague constraint more precise by feeding in some of our *theory of meanings*. Yet *that* degree of precision seems undesirable in an initial attempt to specify the subject matter.<sup>9</sup>

This completes my explication of the basic task: “being a meaning” is a property of certain properties, the ones which play a semantic role and hence which it serves our semantic purposes to ascribe. We have avoided the *ad hoc* charge.

This explication relates the basic task closely to another that I shall call the *normative task*. This is the task of explaining the natures of the properties we *ought to ascribe* for semantic purposes. Clearly, for an ascription to serve those purposes, the property ascribed must have one or both of the semantic roles. So, according to the explication, the property we ought to ascribe must be a meaning. On the other hand, *prima facie*, if something is a meaning then we ought to ascribe it for semantic purposes. Perhaps the reasons for thinking that a property plays a semantic role and hence counting it a meaning are not always sufficient to make it something we should ascribe for semantic purposes. If the property is really not one we ought to ascribe, however, the question of whether or not to count it a meaning becomes uninteresting: *interesting* meanings are ones we ought to ascribe. So, it will do no harm to adopt the *prima facie* view. So, we can assume that a property is a meaning if and only if it is one we ought to ascribe for semantic purposes.

<sup>9</sup> It may be an advantage of the liberalism that it makes eliminativism about meaning harder: in general, the less that is essential to “being an *F*” the harder it is to show that there are not any *F*s.

In the light of this, my *first methodological proposal* is that we should tackle the basic task by tackling the normative one.

We are already familiar with tasks of saying what we ought to ascribe for certain purposes. So one advantage of this methodological proposal is that it relates the basic task to these other ones. Another advantage is the obvious contrast between the normative task and what I will call the *descriptive task*. This is the task of explaining the natures of the properties we do ascribe in attitude ascriptions for semantic purposes; the task of explaining the semantic status quo. It is what most people working in semantics — philosophers, linguists, and psychologists — are in effect doing. Yet it is very different from the normative and basic tasks. Note particularly that the properties it investigates are *meanings* only if ascribing them really does serve our semantic purposes; only if they really play semantic roles. Once the descriptive task is sharply distinguished from the others, the question of its bearing on them arises. I shall consider this bearing in section VIII.

#### VI. OTHER VIEWS OF SEMANTIC TASKS

Semantics very often proceeds without *any* attempt to explicate the talk of “meaning” used, implicitly or explicitly, to define the task(s): the talk is simply taken for granted. Given the *prima facie* unclarity and vagueness of the talk, this practice is surely unacceptable: it leaves us with no firm idea of what the issue is and with the likelihood of being at cross purposes. Once an explication is offered, there is little point in a verbal dispute about the appropriateness of using ‘meaning’ and hence ‘semantics’ in the way proposed. There is a lot of point, however, in asking why a task defined in terms of ‘meaning’ thus explicated is worthwhile. We are, of course, free to study anything. But if semantics is a genuine science, as the naturalist thinks it must be, we should be able to say why it is interesting. A definition that does not say will be *ad hoc*, an appropriate target of Lycan’s mockery. Indeed, a definition of a semantic task in *any* terms which does not demonstrate the worth of the task is an appropriate target.<sup>10</sup>

My discussion of the semantic tasks is an attempt to meet these demands for explication and worth. How do other views of the task fare? This is not the place for a detailed answer, but some brief remarks are in order.

<sup>10</sup> Similar remarks apply to psychological research into the nature of concepts for this is, in effect, investigating the meanings of mental tokens: we need an explication of the talk of concepts together with a motivation for the research arising from the role in psychology of concepts thus explicated.

(1) It is common to define the semantic task in terms of "truth" and "reference." I think that truth and reference are indeed central to *accomplishing* the semantic task but have no place in the *definition* of it. If the definition is not to be *ad hoc*, we need to know why we should be interested in truth and reference. Implicitly, if not explicitly, the motivation is usually thought to come from the role of truth and reference in explaining meaning. But then the definition of the task ought to be in terms of meaning in the first place; truth and reference are not basic enough.

The point is illustrated by the range of apparently semantic disputes that could be arbitrarily settled by a definition in terms of truth and reference:

(i) Some theorists<sup>11</sup> think that truth and reference are deflationary and so think that they have nothing to do with explaining meaning. On this view, to say that, in the specified circumstances, a sentence 'S' is true, is just to say that, in those circumstances, S. And to say that, in the specified circumstances, a name 'a' refers to *x*, is just to say that, in those circumstances, *a* is *x*. So views about the truth conditions and reference are really about extra-linguistic reality and are irrelevant to semantics. Deflationists have in mind an explanation of meaning in other terms altogether, usually some sort of verificationist or "use" theory. Clearly, deflationism should be confronted with arguments, not ruled out by a definition.

(ii) Among those who agree that the properties of having certain truth conditions and of referring to certain objects constitute meanings, there is disagreement about whether such properties *exhaust* meanings. To adopt a definition of the semantic task that favors the view that these properties do exhaust meaning, as direct-reference theorists sometimes seem to, is to settle this dispute by fiat.

(iii) Finally, among those who think that the properties of having certain truth conditions and of referring to certain objects do not exhaust meanings, there is disagreement about what must be added. Fregeans think that we must add a *mode* of presenting a situation or object. Two-factor theorists think that we must add a nonrepresentational conceptual-role factor. To adopt a definition in terms of truth

<sup>11</sup> See particularly, Robert Brandom, "Reference Explained Away," this JOURNAL, LXXXI, 9 (September 1984): 469-92, and "Pragmatism, Phenomenalism, and Truth Talk," *Midwest Studies in Philosophy, Volume XII: Realism and Antirealism* (Minneapolis, Minnesota UP, 1988), pp. 75-94; Paul Horwich, *Truth* (Cambridge: Blackwell, 1990). My "Minimalist Truth: A Critical Notice of Paul Horwich's *Truth*," *Mind and Language*, vi (1991): 273-83, explores the significance of deflationary truth.

and reference which favors the Fregean view is to beg the question against the two-factor theorist.

To settle these disputes, we need to see what role truth and reference do play in the explanation of meaning.<sup>12</sup> And to do that we need an independent account of the task, if not the one I have given, then some other one.

(2) Many think that the semantic task is to explain our "linguistic competence" or "understanding," our "grasp of" meanings. This task is related to that of explaining meanings and can surely be shown to be worthwhile. So, provided it is not conflated with that of explaining meanings, which it usually seems to be,<sup>13</sup> I can have no objection to it. I would argue, however, that the theory of the meanings we grasp is, in an important sense, prior to the theory of our grasp of them.

(3) My tasks concern tokens. There are various tasks concerning types: for example, to describe the conventional meanings of sounds, inscriptions, etc., in a certain language, as dictionaries do; to describe and explain changes in such meanings, as some linguists do; to explain the nature of these conventional meanings in terms of, say, regularities in speaker meanings, as some philosophers do. All of these tasks are clearly worthwhile but none of them is, I think, semantically fundamental.

(4) Some linguists<sup>14</sup> see the semantic task as that of explaining such matters as "synonymy," "similarity of meaning," "meaning redundancy," "ambiguity," and "entailment." Why should we be interested in these matters? The implicit answer is that these are matters about which the speakers have semantic intuitions. But this is unsatisfactory for the reasons already indicated. A better answer would acknowledge that these explanations are contributions to the broader task of explaining *meaning* and so need supplementation by more basic explanations of meaning.<sup>15</sup> The answer would then need

<sup>12</sup> So Ernest Lepore and Jerry Fodor, "Reply to Critics," *Philosophy and Phenomenological Research*, LII (1993): 673–82, are wrong in claiming that I think that it is "a truism that whatever determines reference is ipso facto meaning" (p. 674). I think that it is a substantive truth requiring a great deal of argument.

<sup>13</sup> For example, Dummett, "What is a Theory of Meaning?" in *Mind and Language*, S. Guttenplan, ed. (New York: Oxford, 1975), claims, "A theory of meaning is a theory of understanding" (p. 99); cf. my *Designation*, pp. 92–5. Similarly, linguists seem to conflate explaining syntax with explaining syntactic competence; Kim Sterelny and my, "Linguistics: What's Wrong with 'the Right View'," *Philosophical Perspectives*, III: *Philosophy of Mind and Action Theory* (1989): 497–531.

<sup>14</sup> For example, Jerry Katz, *Semantic Theory* (New York: Harper and Row, 1972), pp. 5–6.

<sup>15</sup> Without the supplementation the explanations rest on what David Lewis, "General Semantics," in *Semantics of Natural Language*, D. Davidson and G. Harman, eds. (Boston: Reidel, 1972), has nicely called, "translations into Markerese" (p. 18).

to go on to explicate this talk of meaning in a way that motivates our interest in it.

(5) Insofar as “formal semantics” is concerned simply with the model-theoretic study of abstract formal languages its task is, of course, quite different from those I have described. Insofar as it is claimed to contribute to the task of explaining meanings in natural languages, it needs to be accompanied by an independent account of that task in order for us to assess the claim.<sup>16</sup>

(6) Davidsonians take the semantic task to be that of saying how to construct a theory of “radical interpretation” for a language. From my perspective, this task is worthwhile but not fundamental enough: it rests on semantic notions that need independent explanation. Davidsonians would disagree, for they think that there is no more to these notions than would be revealed by accomplishing their task. This reflects a very different view of the mind from mine. I favor a robust mentalism that takes thoughts to be objective states posited independently of language (section III). Thoughts are in an important respect “prior” to the language that expresses them and can be used in the explanation of linguistic meaning. Davidsonians, influenced by Quine, start the explanation of language from a behaviorist assumption: “meaning is entirely determined by observable behavior, even readily observable behavior.”<sup>17</sup> This leads to the thesis of the indeterminacy of translation and to a generally antirealist view of meaning. Meanings are not for the most part objective properties with natures awaiting our discovery, and they do not play

<sup>16</sup> Katz, “An Outline of a Platonist Grammar,” in *Talking Minds: The Study of Language in Cognitive Sciences*, T. Bever, J. Carroll, and L. Miller, eds. (Cambridge: MIT, 1984), pp. 17-48 (reprinted in *The Philosophy of Linguistics*), thinks that the central issue in the philosophy of linguistics is the choice between three views of what linguistics is about: physical sounds — the “nominalism” of Leonard Bloomfield; psychological states — the “conceptualism” of Noam Chomsky; or abstract objects — the “Platonism” that Katz himself now urges. This trichotomy does not sit well with the view of the task I have presented. I take semantics to be about physical objects. Sometimes these are sounds, sometimes inscriptions, sometimes radio waves, and so on. But, importantly, sometimes they are thoughts and hence also psychological states. My view does not accept the theoretical constraints — antimentalism and anti-abstractness — of Bloomfield’s, yet it does not follow Chomsky in taking linguistics to be about psychological states of competence. It takes all the objects that linguistics is about to be concrete tokens and so to that extent it is nominalistic. Where it stands ultimately on the nominalism issue depends, of course, on what we make of its ascription of meaning *properties* to those objects. It seems unlikely, however, that the nominalist would have any *special* difficulty paraphrasing away this property talk.

<sup>17</sup> Davidson, “The Structure and Content of Truth,” this JOURNAL, LXXXVII, 6 (June 1990): 279-328, here p. 314; cf. Quine, “Two Dogmas Revisited,” *Canadian Journal of Philosophy*, XXI (1991): 279-328, here p. 272.

roles in explanatory theories of mind and language. The only reality captured by meaning talk is a set of verbal dispositions.<sup>18</sup> Beyond that there is nothing but our practice of interpreting each other using principles of charity and rationality, a practice which should be seen more as *imposing* a semantic reality than as discovering one.<sup>19</sup> This "interpretative perspective," reminiscent of the European *Verstehen* tradition,<sup>20</sup> has no place for the sort of tasks I have defined.<sup>21</sup>

The Davidsonians have given no argument for the behaviorism that underlies their perspective. From my perspective, that behaviorism seems like just another dogma of empiricism.<sup>22</sup>

In conclusion, I emphasize that I am *not* claiming that the properties that I have called "meanings" are the *only real* meanings nor that the tasks I have called "semantic" are the *only proper* tasks for semantics. I doubt that there is any interesting matter of fact about such claims. I do claim that those properties are worth investigating and that those tasks are worth performing. I claim further that those tasks are appropriately fundamental. Perhaps all this is true of other tasks, but that always needs to be demonstrated.

This completes my discussion of the semantic tasks. One further matter must be discussed before making any further methodological proposals.

#### VII. DOES A TOKEN HAVE MORE THAN ONE MEANING?

Does a token have more than one meaning? Note that this question does not concern the familiar matter of ambiguity, which is a property of *types*.

Ordinary talk of "*the meaning*" of a word encourages the view that a token *does not* have, indeed *cannot* have, more than one meaning. This view is taken over in our semantic theories and disputes; for example, the dispute over semantic holism seems to be over whether

<sup>18</sup> Quine, *Word and Object* (Cambridge: MIT, 1960); "Mind and Verbal Dispositions," in *Mind and Language*, pp. 83-95.

<sup>19</sup> Davidson, *Inquiries into Truth and Interpretation* and "The Structure and Content of Truth."

<sup>20</sup> G. MacDonald and P. Pettit, *Semantics and Social Science* (New York: Routledge, 1981), have derived a *Verstehen* view from an explicitly Davidsonian perspective. They have also claimed that this is the only good route to *Verstehen*.

<sup>21</sup> It should be noted that some take over central ideas from Davidsonian semantics without adopting his interpretative perspective; for example, Lycan, *Logical Form in Natural Language*, R. Larson and P. Ludlow, "Interpreted Logical Forms," *Synthese*, xcvi (1993): 305-57.

<sup>22</sup> On this, see Louise Antony, "Naturalized Epistemology and the Study of Language," in *Naturalistic Epistemology*, A. Shimony and D. Nails, eds. (Boston: Reidel, 1987), pp. 235-57. I have discussed Davidsonian semantics elsewhere: *Designation*, pp. 115-8; *Realism and Truth*, pp. 186-99.

the one and only meaning a token has is holistic or localistic. Yet the view seems to be nothing but a prejudice.<sup>23</sup>

I have noted that meanings are usually thought to be constituted by inferential relations among tokens and/or direct causal relations to the world. A token has many such relational properties and clearly different sets of them could constitute indefinitely many *candidates* to be meanings. To tell whether one of these really is a meaning we have to see whether it has one of the semantic roles. Prior to arguments one way or the other on this, we ought to take the property's candidature seriously. And we ought to take the possibility of a token having more than one meaning very seriously.

Given the situation in other areas, we should expect to find that a token does indeed have more than one. Thus, a particular person may have a variety of economic properties: "being a capitalist," "a landowner," "a banker," and so on. Each of these properties plays a role that economic properties are supposed to: for example, explaining economic behavior. We should expect that a particular token in thought or language will have a similar variety of semantic properties. Each of these will play a semantic role: one might explain one bit of behavior, another, another, and a third might serve as a guide to external reality.<sup>24</sup>

If the early Quine is anywhere near right in his view of ordinary attitude ascriptions, then the folk ascribe more than one property to a token for semantic purposes. Quine distinguishes between "transparent" ascriptions, where only the reference of a token is of interest, and "opaque" ascriptions, where some finer-grained meaning of the token is of interest.<sup>25</sup>

<sup>23</sup> There is one respect in which some theories do not seem to follow the prejudice: under the influence of H. P. Grice they allow that a token has both a speaker meaning and a literal or conventional meaning, although these meanings are usually the same. What theories do not contemplate is that a token might have more than one speaker meaning, or more than one conventional meaning. That is the prejudice which concerns me here.

<sup>24</sup> The view that a token has more than one meaning seems to be an idea whose time has come. It has been independently suggested by Akeel Bilgrami, *Belief and Meaning: The Unity and Locality of Mental Content* (Cambridge: Blackwell, 1992), and Eric Lormand, "A Marriage of Holism and Atomism" (unpublished).

<sup>25</sup> And I suspect that the folk would be prepared to call either of these properties "meanings." If not, this would be an example of the way that the ordinary use of 'meaning' was not perfectly suited to theoretical semantics. What we are considering here is the ordinary application of 'meaning' to *tokens*. It is probably more often applied to *types*, to what convention makes common to tokens, the sort of thing that is to be found in a dictionary. So, the meaning of a type that has tokens that do not conventionally have the same referent — for example, the type 'she' — is not thought to include a referent. In contrast, if the tokens do conventionally have the same referent — as, for example, tokens of 'echidna' do — ordinary usage may often treat that referent as the meaning of the type.



Consider also the popular two-factor theories of meaning. These theories often describe purposes for ascribing meanings which are similar to the two I have outlined. They then assign two relatively independent meaning factors to a token to serve those respective purposes, a conceptual-role factor to explain behavior and a truth-referential factor to serve as a guide to reality. They prefer to say that these factors jointly constitute *the* meaning of a token rather than that each severally constitutes *a* meaning of the token, but this preference seems to reflect only the prejudice that a token must have just one meaning.

In sum, I think that we should leave open the possibility that we ascribe more than one property to a token for semantic purposes (descriptive), that we ought to ascribe more than one (normative), and hence that a token has more than one meaning (basic).

#### VIII. THE METHODOLOGY FOR THE NORMATIVE/BASIC TASK

My first methodological proposal was that we should tackle the basic task by tackling the normative one. I turn now to the methodology of that task.

What should be the basis for our choosing to ascribe one rather than another of the many properties of a token that are candidates to be meanings? The answer is implicit in the task. We should ascribe the ones which play semantic roles and hence which it serves our purposes to ascribe: the ones which play the role outlined in explaining behavior and informing us about reality.

We know from the history of science that it is not, in general, easy to tell which properties we should ascribe to serve our purposes. We know from the philosophy of science that it is not easy to say how to tell. The first sign of the modesty of my aims is that I shall attempt to say only a little on how we should tell what to ascribe in semantics.

I think, however, that the little I shall say is important. It concerns the bearing on this normative task of the descriptive semantic task, the task of explaining the natures of the properties we do as a matter of fact ascribe for semantic purposes. Perhaps the most unfortunate consequence of unclarity about the semantic tasks has been the lack of attention to the relations between the normative and the descriptive. On the one hand, semanticists seem to be concerned, most of the time, solely with the descriptive task. Yet, *meanings* should be our primary concern in semantics, and meanings are what we *ought* to ascribe for semantic purposes *not necessarily* what we *do* ascribe. So the descriptive task should proceed with at least an eye on the normative one. On the other hand, when normative questions are raised they are usually addressed without close attention to the descriptive

facts. In this section, I shall argue that these facts are very relevant to the normative task.

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 for a basic one.

We should go further: given the apparent success of our current practices, it is unlikely that any normative proposal uninformed by a descriptive view will be plausible. Consider, for example, the epistemic problems faced by a revisionist proposal. How could we tell that nothing has the properties we currently ascribe to thoughts and utterances for semantic purposes without attention to the natures of those properties? How could we tell that those properties fail to play their semantic roles without a good deal of knowledge of what they do do? How could we assess the claim that some "new" properties do a better job than those "old" ones, including explaining the apparent success of the old, without having a clear picture of the old? 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.

In the light of this discussion, I make my *second methodological proposal*: we should look to the descriptive task for evidence for the normative/basic one. Further, I suggest, although I shall not argue, that the main, if not the only, justification for the usual focus on the descriptive task is the bearing of that task on the normative one.

In the next section, I turn to the methodology of descriptive tasks in general. I shall consider this with an eye to its application to semantics. The application is in the following section.

#### IX. THE METHODOLOGY FOR DESCRIPTIVE TASKS IN GENERAL

What constitutes some property we ascribe, "being an *F*"? The question may often have little scientific interest, of 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 *F*s “in all possible worlds?” What is the “ultimate” method?

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*.” My focus will be on the preliminary first stage, for it is here that there is a place for the sort of intuitions and thought experiments that are so prominent in semantics. But I start with the second stage.

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. Modesty will again prevent my trying to say (including trying to say what role, if any, intuition plays in accomplishing 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 *F*s are not; for example, whales are not fish. We may conclude that some things identified as non-*F*s are *F*s; for example, tomatoes are fruits. We may even conclude that *none* of the things identified as *F*s 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.

Turn now to the preliminary first stage. How do we identify the examples that are to be examined in the second stage? Sometimes, once again, we have a well-established theory to help with the job; for example, Mendelian genetics helps identify the genes that are to be examined by molecular genetics. But suppose that we do not, which is, once again, surely the right supposition in semantics. In the absence of reliable theory, we must consult the experts about *F*s and see what *they* identify as *F*s and non-*F*s. Often these experts will be scientists. Sometimes they will be engineers, tradespeople, and the like. And sometimes they may be just plain folk. But, whoever they are, our first stage consists in conducting “identification experiments” on them.

What we elicit in these experiments are the expert’s most basic “intuitions” about “being an *F*.” We can elicit other intuitions, asking the expert not only to identify *F*s but also to tell us about them.

These other intuitions are clearly richer and less basic than the identification ones: a person may be good at recognizing *F*s without having much reliable to say about them; this may be the situation of the folk with pains.

When we are starting from scratch, we need the identification intuitions but we do not need the richer ones. This is not to say that we should not use them. They may well be a useful guide to what we shall discover at the second stage.

The generalizations of these intuitions constitute a theory. There may be no reason to suppose that the experts consulted *represent* this theory and so we should be careful about attributing the theory to them. Nevertheless, suppose, as is quite likely, that they would readily assent to the theory if it were presented to them. It is appropriate then to say that the theory is their "tacit" theory. If the experts are the folk, as they presumably are with pains for example, then the tacit theory is "folk theory."

So, we have found a place for intuitions. It is important not to exaggerate that place. At best, the intuitions are likely to be seriously incomplete, reflecting only part of the theory we need. And they may be wrong: they are empirical responses to the phenomena. Our earlier brief consideration of the second stage shows that even the most basic intuitions, expressed in identification experiments, are subject to revision in the face of scientific examination.

Intuitions are often needed to identify the subject matter for the descriptive task, and may be otherwise helpful, but nothing ultimately rests on them.

A worry may obtrude about identification experiments. "The identification of *F*s will consist in *calling* things "*F*." So you are assuming that if people who both understand '*F*' and have the expertise to recognize things that have the property, "being an *F*," call something "*F*," then it is likely to have that property (if anything has). This seems to be semantic. Yet you plan to apply conclusions drawn from this discussion to semantics." The "identification assumption" is indeed semantic but it is so minimally so that it should be acceptable to semanticists of all persuasions. Its main interest is epistemological.

We have found a place for intuitions, but not yet for thought experiments. They can have a role in identification. Instead of "real" experiments that confront the expert with phenomena and ask her whether they are *F*s, we confront her with *descriptions* of phenomena and ask her whether she *would say* that they were *F*s. These thought experiments provide valuable clues as to what the expert would identify as an *F* and a non-*F*. They can do more: the descriptions that elic-

it the response indicate the richer intuitions that we have already noted can be a useful guide to the second stage. Some experiments may be difficult, perhaps impossible, to perform other than in thought.<sup>26</sup> Valuable and useful as thought experiments may be in practice, they are dispensable in principle: we can make do with real experiments. And the results of thought experiments have the same empirical status as the results of real experiments.

With an eye to our later discussion of language, we should note two special situations. The first arises if the *F*s we are investigating are the products of a human skill or competence; for example, if we are investigating the nature of a good tennis shot or a good chess move. Someone who has the relevant competence has ready access to a great deal of data that are to be identified as *F*s. She does not have to go out and look for data, her competence produces them. As a result, she is in a very good position to become an expert: she simply has to reflect upon her own performance. But note that this is not to say that she *will* become an expert: a person can be competent and yet reflect little on the output of that competence; knowledge-how is not knowledge-that. And even if she does become an expert, her opinions are still subject to revision by the second-stage examination. She is privileged in *her ready access to data*, not in *the conclusions she draws from the data*. Finally, to say that competence may lead to expertise is not to say that the incompetent cannot be experts, perhaps even better experts.

The second special situation arises if theorists count themselves among the identification experts. They then have the advantage of not having to consult other people to conduct real or thought experiments.

Thought experiments performed in this second situation are the characteristic "armchair" method of philosophy, a famous example of which is "the analysis of knowledge." So what I am proposing here amounts to a naturalistic account of that method.<sup>27</sup>

In the last section, I argued that the descriptive task bears on the normative. Similarly, the normative bears on the descriptive, yielding another source of evidence. Suppose our current practice of ascribing certain properties is successful. Then it is likely that any properties that we ought to ascribe will include those ones. So any *independent* evidence we can get about the nature of the properties that we

<sup>26</sup> As David Chalmers pointed out to me.

<sup>27</sup> The naturalist does not deny "armchair" intuitions a role in philosophy but denies that their role has to be seen as a priori: the intuitions reflect an empirically based expertise at identification; cf. George Bealer, "The Incoherence of Empiricism," *Proceedings of the Aristotelian Society*, Supp. Vol. LXVI (1992): 99-138.

ought to ascribe — what ones it would serve our purposes to ascribe — is evidence about the nature of the properties that we do ascribe. (The evidence for the normative position must be independent, of course, to avoid a circle: it must not depend on the evidence for the descriptive position.)

One final methodological point. Some problems are not in the evidence but in reality: there may simply be no fact of the matter whether or not certain objects are *F*. Consider for example, a case where the various criteria of specieshood come into conflict; or consider the intermediate stages in the evolution of one species from another. The nature of “being an *F*” is, to that extent, indeterminate or — given the semantic associations of that term — perhaps we should say, vague. This raises some interesting logical puzzles<sup>28</sup> but is otherwise an untroubling fact of life.

#### X. THE METHODOLOGY FOR THE DESCRIPTIVE TASK IN SEMANTICS

We now apply these findings to semantics. We must say how we should tell what constitutes the properties we ascribe for semantic purposes. These “*putative* meanings” will *be* meanings if they do indeed play semantic roles.

In the last section, I took a negative view about reliance on semantic theories. Philosophy and linguistics offer a variety of theories, or theory fragments, of meaning. These theories are hardly ever well-enough established, however, to settle the disputes that bedevil semantics. The theories themselves are as much in contention as anything is, part of the problem not the solution. So, I make my *third methodological proposal*: we should use the “ultimate” method on putative meanings to accomplish the descriptive task. And, in applying that method, we should not rely on those semantic theories to identify examples.

So little is established and settled in semantics that the need to use this “ultimate” method is pressing. Yet it is seldom explicitly applied to putative “meanings.” What we find in philosophy is that it is, in effect, frequently applied to “truth conditions” and “reference.” And what we find in linguistics is that it is, in effect, frequently applied to “synonymy,” “similarity of meaning,” and the like. But these are not the right places to apply the method to settle some of the most burning issues in semantics: they are not basic enough (section VI).<sup>29</sup>

<sup>28</sup> See, for example, Francis Jeffrey Pelletier, “Another Argument against Vague Objects,” this JOURNAL, LXXXVI, 9 (September 1989): 481-92.

<sup>29</sup> Apart from this, there are two reasons for being a little skeptical about folk identifications of cases of synonymy, similarity of meaning, and the like. First, we should wonder whether the “meanings” involved in these identifications are the

To apply the “ultimate” method to a putative meaning, we must first identify some apparently uncontroversial examples of tokens with that property and of tokens without it; and second we must examine the examples to determine the nature of the property. Who are the experts to be consulted in the first-stage identifications? The folk, in their frequent use of *t*-clauses in attitude ascriptions, are as expert as anyone. And the folk are in an advantageous position to become experts because they themselves are competent to produce token thoughts and utterances to which the properties are ascribed. Finally, theorists can count themselves among the expert folk. So, in the first stage, semantic theorists have two advantages: no need to consult others and ready access to data.

To identify examples in the preliminary first stage, we look to ordinary attitude ascriptions, sentences that ascribe putative meanings to thoughts and utterances using *t*-clauses.<sup>30</sup> We bring these ascriptions to bear on the descriptive problem with the help of “the identification assumption”: if people who both understand ‘*F*’ and have the expertise to recognize things that have the property, “being an *F*,” call something “*F*,” then it is likely to have that property (if anything has).

An illustration of the use of this assumption outside semantics will pave the way for its use in semantics. Consider the role that the assumption might once have played in the case for a theory of, say, echidnas. To discover what “being an echidna” amounts to, we must

---

same as the meanings we are investigating. Lycan’s double indexical theory of meaning probably applies to the folk as much as to theorists. Second, these identifications involve terms like ‘synonymy’ drawn from semantic theory and are rather distant from ordinary concerns. In contrast, identifications of tokens with a certain putative meaning need involve no such terms and are central to the folk’s ubiquitous practice of ascribing thoughts and utterances. So we should expect the folk to be less expert at identifying synonymy and the like than putative meanings.

The second of these reasons for doubt applies also to folk identifications of *ungrammatical* tokens. ‘Ungrammatical’ is a term from syntactic theory and folk identifications of what it applies to, unlike of, say, “odd things for someone to say,” are probably not common. The intuitions about grammaticality which play a key role in the study of syntax reflect a good deal of theory beyond folk theory.

<sup>30</sup> We can apply the “ultimate” method immediately to these putative meanings because the folk are experts at identifying them, as the folk’s use of *t*-clauses reveals. We could not have applied the method earlier to discover the nature of the second level property, “being a meaning,” precisely because nobody has the required expertise (section II). When we start semantics, we have access to apparently uncontroversial identifications of some “determinates” that are putative particular meanings but not of the “determinable” that these properties are candidates to share, “being a meaning.” It was this lack that led to my attempt to specify a subject matter worthy of semantic investigation (sections III-V). In semantics, we really start from scratch.

first use experts to find some uncontroversial examples of echidnas and nonechidnas. Applying the assumption, we discover the examples: if people who both understand 'echidna' and have the expertise to recognize things that have the property, "being an echidna," call something "echidna" then it is likely to have that property (if anything has).

We do the same in semantics for a theory of, say, the putative meaning, "this is an echidna." To discover what having that property amounts to, we must first use experts to find some uncontroversial examples of tokens with the property and of tokens without it. Applying the assumption, we discover the examples: if people who both understand 'that this is an echidna' and have the expertise to recognize tokens that have the property, "this is an echidna," call something "that this is an echidna," then it is likely to have that property (if anything has).<sup>31</sup>

This experiment elicits the expert's most basic intuitions about the property, "this is an echidna," reflections of her tacit theory. We could probe that theory further, revealing richer intuitions about the circumstances in which tokens with that putative meaning would be true, about what parts of them refer to, about what would verify them, and so on. These richer intuitions are a guide to the second stage.

It is important to note that my application of the term 'tacit theory' here is quite compatible with my earlier rejection of its application to our conceptual competence (section I). The present application is to generalizations of explicit intuitions about particular putative meanings.<sup>32</sup> The basis for rejecting the earlier one was precisely that our conceptual competence — a capacity to think certain thoughts — does *not* consist in such semantic intuitions. The relation of competence to tacit theory is explored further below.

<sup>31</sup> We have to keep in mind that 'that this is an echidna' may be ambiguous, specifying more than one meaning (section VII).

<sup>32</sup> "Simulation theory" would reject this application: R. Gordon, "Folk Psychology as Simulation," *Mind and Language*, I (1986): 158-71; Alvin Goldman, "Interpretation Psychologized," *Mind and Language*, IV (1989): 161-85. I have no quarrel with the theory's denial that these generalizations are *represented* in us. I take it, however, that the theory goes further, denying that the generalizations are "psychologically real" at all: even that we embody them without representing them in the same way that a simple calculator embodies the laws of addition without representing them; that they have become "hard-wired" into us. If simulation theory were right about this, then my present talk of "tacit theory" would be inappropriate. But I follow Stephen Stich and Shaun Nichols, "Folk Psychology: Simulation or Tacit Theory?" *Mind and Language*, VII (1992): 35-71, in finding the theory poorly supported by the evidence. I think that the embodiment of these generalizations provides the best explanation of our having the semantic intuitions. Should I be wrong, the talk could be dropped from my argument without cost.



Thought experiments can have a prominent role. We describe situations and ask whether we, and other folk experts, would apply the *t*-clause 'that this is an echidna' to them. These thought experiments provide valuable clues as to what the expert would identify as tokens with the property, "this is an echidna," and as tokens without it and, via the descriptions that elicit the response, evidence of the expert's richer intuitions. Valuable and useful as these thought experiments may be, we could make do with real experiments.

With the results of the preliminary first stage in hand, we move to the second stage: we examine the alleged examples to determine the nature of the putative meaning, "this is an echidna." This stage is both hard to do and hard to say how to do. We can say that, however, just as this stage led us to a theory powerful enough to overturn intuitions about what things are fish, fruit, and witches, it might lead us to one that overturns intuitions about what tokens have the property, "this is an echidna." Even more, it might lead us to overturn the richer intuitions of folk semantics. And it will surely show them to be seriously incomplete. In some cases, it may give no answer because there is no answer to be had: meanings may be vague like the rest of reality.

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 a 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 a meaning.

This point about intuitions does not need to be modified because of the special situation in semantics: the expert's intuitions may be partly the result of reflecting on the product of her own competence. As a result of this competence, the expert herself produces tokens that may refer to *x*, be true in certain circumstances, be verified in such and such a way, and so on. So she is privileged in having *ready access* to such data. But this does not give any privilege to *the conclusions that she draws from the data*; for example, that they *do* refer to *x*, *are* true in those circumstances, or *are* so verified. These conclusions of the competent, just like those of the incompetent, are empirical responses to the phenomena and open to question.<sup>33</sup>

<sup>33</sup> Similar remarks apply to linguistic intuitions. So I think that the small amount of sympathy I showed elsewhere ("Linguistics: What's Wrong with 'The Right View'," p. 523) to the view that these intuitions are "the voice of competence" is a mistake: the intuitions are the result of central processor reflection upon tokens of the sort that are and are not the product of syntactic competence.

All this having been said, we should not expect our semantic intuitions to be overturned. I think that the considerations already adduced, and one to come in the next section, give good reason for confidence in these intuitions, so far as they go.

Finally, any independent evidence that we can get on the normative issue of the meanings we ought to ascribe is also evidence on the descriptive issue of what meanings we do ascribe. For, given the apparent success of our current ascriptions at serving our semantic purposes, it is likely that some of the ones we ought to ascribe will be ones we do ascribe.

Before concluding, we need to look briefly at the meanings of attitude ascriptions themselves.

#### XI. THE MEANINGS OF ATTITUDE ASCRIPTIONS

There is undoubtedly something confusingly different about the use of the identification assumption in semantics because of the unusual way in which we ascribe putative meanings. To ascribe any sort of property we must, of course, use an expression for that property; thus, to ascribe "being a gene" we use 'gene' and "being an atom" we use 'atom'. These words hardly ever *have* the property they are used to ascribe and if they do it is not in virtue of having the property that they ascribe it; thus 'gene' is not a gene and although 'short' is short it is not because it is short that it ascribes shortness. Yet, in the case of meanings, the situation seems different. The expression 'that *p*' used to ascribe a putative meaning often seems to ascribe that property in virtue of the fact that its content sentence, '*p*' has the property (or one very like it). So there always seems to be "an intimate link" between the two properties. In applying the "ultimate" method, however, we can abstract from this link, treating the *t*-clause 'that *p*' like any arbitrary expression for a property.

I have emphasized that *competence with* concepts/language is an ability that is logically distinct from any *expertise about* concepts/language. Yet the "intimate link" suggests an important causal relation between them. Perhaps expertise at applying 'that *p*' to the tokens of others is an essential part of the usual causal process of becoming competent with '*p*'. So a human being in a normal situation would not acquire that competence without becoming an expert at identifying tokens with the property, "*p*." Competence is *not* a tacit theory, but perhaps the theory *would normally accompany* the competence. This would add to the prima facie case for folk expertise at identification.<sup>34</sup>

<sup>34</sup> I am indebted to Carol Slater for a criticism that prompted this paragraph.

I have said that the “ultimate” method is seldom explicitly applied to meanings. Yet, it may be objected, the close attention to attitude ascriptions recommended by this application is surely common. This attention is indeed common but its focus is different. The method recommends a focus on the ascriptions as evidence of *the putative meanings ascribed*, “first-level” meanings. The common focus is on the ascriptions as evidence of *their own meanings*, “second-level” meanings.

This second-level investigation is almost always an application of the “ultimate” method to the truth conditions of ascriptions: we consult our intuitions about the situations that make a token ascription true. But, again, to assume that truth conditions are all there is to a sentence’s meaning, or even that they have anything to do with it, is to beg a lot of semantic questions (section VI).

How then should the second-level investigation proceed? We can of course apply the “ultimate” method to putative meanings here too. This will require attention to ascriptions of attitudes *to attitudes*. And perhaps we can avoid the charge of question begging in applying the method to truth conditions if we can first establish the plausibility of truth-referential semantics at the first level. Finally, we may hope to trade on the intimate link: evidence of the meaning ascribed is also evidence of the ascription’s meanings.

#### XII. CONCLUSIONS AND PREDICTIONS

From a naturalistic perspective, contemporary semantics has some troubling features: the resort to intuitions and thought experiments; the inexplicit reliance on apparently *ad hoc* accounts of the semantic tasks; broadly, the lack of a scientifically appealing method for settling the disputes that bedevil it. I have attempted to remedy this. I have defined the tasks by looking to the purposes we attempt to serve in ascribing meanings, and have proposed a methodology for accomplishing the tasks. We should tackle the “basic” task of explaining the nature of meanings by tackling the “normative” one of explaining the properties that we ought to ascribe for semantic purposes (“first proposal”). Our ordinary attitude ascriptions ascribe certain properties for semantic purposes. These properties are *putative* meanings. Given the apparent success of the ascriptions it is likely that these putative meanings are *real* ones. So we should look to the “descriptive” task of explaining putative meanings for evidence for the normative/basic one (“second proposal”). Since we approach the descriptive task pretty much from scratch, we should use the “ultimate” method (“third proposal”). The preliminary first stage of this method identifies examples for a straightforwardly scientific examination in the second stage. Intuitions and thought

experiments of the sort that dominate semantics are important in the first stage. They are empirical responses to the phenomena, however, and are open to revision at the second stage.

I predict that applying this methodology will help with all semantic issues and will quickly settle some, including some of the most notorious.<sup>35</sup>

(1) We shall discover that we do indeed ascribe more than one putative meaning to a token; indeed, I think that we ascribe as many as four to “self-demonstratives” like ‘I’. Not only do we ascribe more than one, we ought to, for our semantic purposes are served by so doing. So tokens have more than one meaning.

(2) The evidence against semantic holism as a descriptive doctrine will be decisive. Consider any putative meaning. We shall find that the folk are prepared to ascribe it to tokens that differ enormously in their inferential properties. So, few if any inferential properties could constitute the ascribed property; it must be localistic. Given the success of folk ascriptions in serving our semantic purposes, we ought to ascribe such localistic properties. We shall find independent support for this normative conclusion from the fact that only localistic properties have the sort of generality in which we are interested. So holism is wrong not only about putative meanings but about meanings.

(3) The direct-reference theory of proper names will turn out to be partly right and partly wrong. It is right in claiming that one of the properties we do, and should, ascribe to a name is simply its role of referring to its bearer. It is wrong in insisting, so determinedly, that this is *the* (one and only) meaning of a name. The evidence will show that we often ascribe a property constituted by a certain mode of referring to the bearer. And we ascribe this property as well because doing so serves our semantic purposes. Direct-reference theorists have no principled basis for denying this finer-grained meaning. Their position is theoretically *ad hoc*.

(4) We shall find that there is indeed an intimate link: the meaning ascribed is always identical or closely related to one of the meanings of the content sentence in the *t*-clause. And it is in virtue of this

<sup>35</sup> I attempt to confirm these predictions in *Coming to Our Senses: A Naturalistic Program for Semantic Localism* (New York: Cambridge, forthcoming). The methodology also has consequences for the direction of psychological research on concepts. It suggests the strategy of testing when subjects *will ascribe* a particular meaning, truth value, reference, etc., to a thought or utterance. Thus, rather than the present practice of testing reactions to birds or representations of birds, we test when subjects will ascribe the concept “bird” to a person in an attitude ascription (for example, “x believes that ... bird ...”).

link that the *t*-clause has *its* meanings, meanings arising from its role of ascribing what it does.

(5) I think that the evidence will show that the properties we ascribe for semantic purposes are all truth-referential. So the evidence will count against any descriptive theory that claims otherwise; for example, against two-factor, use, or verificationist theories. Of course, such a theory might still be normatively correct and hence correct about meanings. But it is going to be hard to show this. Why do our ordinary ascriptions *seem* to serve our purposes so successfully if they do not really do so? What reason have we for thinking that the ascriptions recommended by the theory will better serve our purposes? The onus of answering these questions is heavy.

Perhaps another naturalistic methodology reflecting a different view of the tasks would yield different predictions. But any such methodology needs to be set out and justified.

MICHAEL DEWITT

University of Maryland/College Park