Introduction

Three important questions get insufficient attention in semantics. What are the semantic tasks? Why are they worthwhile? How should we accomplish them? The central purpose of this book is to answer these "methodological" questions and to see what semantic program follows from the answers.

It is troubling that much semantic theorizing proceeds with inexplicit reliance on apparently ad hoc views of the semantic tasks. Thus it is common to take for granted that semantics is concerned with truth and reference. I think that this view is right, but *why* is it right? What can we say to someone who disagrees, claiming that semantics should be concerned with, say, warranted assertability or "use"? Furthermore, it is troubling that, in attempting to accomplish the semantic task, we all go in for "intuition mongering," even those of us who are naturalistically inclined and skeptical of the practice (e.g., Jerry Fodor 1990: 169). Broadly, it is troubling that we seem to lack a scientifically appealing method for settling the disputes that bedevil semantics. In Chapter 2, I propose a view of the semantic tasks by looking at the purposes we attempt to serve in ascribing meanings. And I propose a way of accomplishing them. This methodology has a place for intuitions, but it is the same limited place that they have elsewhere in science. I think that applying this methodology will help with all semantic issues. In this book I shall use it in the hope of settling some, including some of the most notorious.

A by-product of this methodological discussion is a naturalistic account of the thought experiments characteristic of "armchair" philosophy.

1

In approaching the methodological questions, I make three
important and related assumptions. First, I assume anti-
Cartesianism. It is common to think that linguistic-conceptual
competence brings "privileged access" to meanings (or contents).
One example of this is the widespread view that semantic compe-
tence consists in knowledge of truth conditions. Another is the
received Fregean view that two expressions that differ in informa-
tiveness must differ in meaning. I argue briefly against such
Cartesianism here (secs. 1.7, 1.8, 2.2) and have argued against it at
much greater length elsewhere (1981a: 95–110; 1983: 674–5;
1991b: 270–5; Devitt and Sterelny 1989). In any case, I think that
the onus lies very much on the other side. The supposition that
someone who has a thought, or uses an expression, that has a
certain meaning *thereby has knowledge about that meaning* is a strong
one requiring much more support than it has ever been given
(even if the knowledge is described as only "tacit"). I think that
we should be skeptical of the supposition that semantic compe-
tence alone yields semantic propositional knowledge. My aim is
for a semantics that does not make these suppositions.

My second assumption is already obvious. It is naturalism: that
there is only one way of knowing, the empirical way that is the
basis of science (whatever that way may be). So I reject "a priori
knowledge." I do not give a detailed argument for my rejection
but I do give two reasons (2.2): Briefly, first, with the recognition
of the holistic nature of confirmation, we lack a strong motivation
for thinking that mathematics and logic are immune from empiri-
cal revision; and, second, the idea of a priori knowledge is deeply
obscure, as the history of failed attempts to explain it shows.

My third assumption is implicit but nonetheless important. It is
realism about the external world: that the physical world posited
by science and common sense objectively exists independently of
the mental. The chances that discoveries about meaning will cast
doubt on this realism are, in my view, just about nil. I have argued
for this at length elsewhere (1991b). I take realism so much for
granted in this work that I hardly mention it.

My methodological discussion was one natural way to start this

book. But I have a particular concern with semantic (or meaning)
holism, and I do not need either the methodology or the previous
three assumptions for my critique of the case *for* this holism. So I
decided to make this critique Chapter 1 and the methodological
discussion Chapter 2.

My aim in Chapter 1 is not to defend an "atomistic" localism
like Fodor's according to which *no* inferential property of a token
constitutes its meaning. It is to defend a more moderate, "molecu-
lar," localism according to which *a few* of the inferential properties
of a token *may* constitute its meaning. And I expect that we shall
discover that many meanings *are* indeed constituted by inferential
properties. In leaving open this possibility, I challenge the conven-
tional wisdom that molecular localism is untenable because there
is "no principled basis" for its distinction among inferential prop-
erties.

My first use of the methodology is in Chapter 3 to argue a case
against semantic holism and *for* molecular localism. Chapter 1
rejects the arguments against there being a principled basis for
distinguishing inferential properties alleged to constitute a token's
meaning from its other inferential properties. Chapter 3 argues
that, insofar as we need a principled basis, we have one.

I use the methodology next, in Chapter 4, to present a program
for a *particular* localistic semantics. This program is "Representa-
tionalist": It holds that the meanings of sentences are entirely
constituted by the properties that go into determining their truth
conditions and that the meanings of words are entirely constituted
by properties that go into determining their references. (So Rep-
resentationalism is in the spirit of the slogan "The meaning of a
sentence is its truth condition.") Arguing for this program requires
rejecting two-factor, functional- (conceptual-) role, verificationist,
and "use" theories. I use the methodology finally, in Chapter 5,
to reject "narrow"-meaning theories and other forms of revision-
ism and eliminativism.

Representationalism is, of course, common in semantics. The
most notable thing about my program is *the argument for it* based
on the proposed methodology. Also notable is the claim that a

token has more than one meaning. With this claim goes partial
acceptance and partial rejection of two influential views of singular
terms: first, the " 'Fido'–Fido," or "Millian," view, recently resur-
rected by direct–reference theorists, that a term's only meaning is
its property of referring to its bearer; second, the Fregean view
that a term's only meaning is its "mode of presenting" its bearer. I
argue that a term has both meanings. I agree with Frege that the
meaning that is a mode *may* be descriptive, involving inferential
links to other terms; that is my molecularism. I disagree with
Frege, and just about everybody else, in arguing that some mean-
ings are nondescriptive causal modes of reference.

I mostly call what I propose a "program" rather than a "theory"
because I do not go into a lot of semantic details; in particular,
although I talk about reference all the time, I say very little to
explain it. This does not reflect any lack of interest in these details
(see my 1981a, for example). Rather, my present aim is to focus
attention on the more general question: Which way should se-
mantics go in future?

Chapters 1 and 3 are on the holism–localism issue. Chapter 3
presupposes the methodological Chapter 2. A reader interested
only in arguments to do with that issue should read only those
three chapters. On the other hand, a reader who does not care
about the holism–localism issue but is interested in other aspects
of the semantic program can skip Chapters 1 and 3 and focus on
Chapters 2, 4, and 5.

SUMMARY OF CHAPTERS

At its most extreme, semantic holism is the doctrine that all of the
inferential properties of a token in language or thought constitute
its meaning. Holism is supported by the consideration that there
is no principled basis for molecular localism's distinction among
these properties. In Chapter 1, I reject four arguments for this
consideration. The first, the argument from confirmation holism,
is dismissed quickly because it rests on verificationism, which the
localist need not accept. The second, the argument from the

rejection of analyticity, is more popular and is discussed at some length. I argue that it fails because it saddles the localist, gratuitously, with epistemic assumptions; in particular, with the Cartesian thesis that if the meaning of a word depends on its inferential relations to other words then a competent speaker must know about this. Localism is a *semantic* doctrine that need not be committed to any particular epistemological thesis. So it need not be committed to a priori knowledge or to knowledge that is in any interesting sense unrevisable. The third is the argument from psychological explanation. I discuss a version of it due to Ned Block, based on Hilary Putnam's "Ruritania" example. I reject the argument because it begs the question. The fourth, the argument from functionalism, needs to be accompanied by a further argument that functionalism is *essentially* holistic. In any case, it could only establish a very mild holism.

In Chapter 2, I address the methodological questions that began this introduction. I define three semantic tasks by focusing on the purposes for which we ascribe meanings: in particular, the purposes of explaining behavior and using thoughts and utterances as guides to reality. I then propose a methodology for accomplishing these 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 attribute 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"). Because 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. However, they are empirical responses to the phenomena and are open to revision at the second stage.

Finally, in doing semantics, we should "put metaphysics first" ("fourth proposal").

In Chapter 3, I present a case for semantic localism. It is generally thought that the molecular localist must show that there is a principled basis for distinguishing any inferential properties of a token that she alleges constitute its meaning from its other inferential properties. I begin the chapter by responding to this demand. We must distinguish two ways of construing it. (*a*) If the demand were making a "descriptive" point, it would require that we distinguish the inferential properties that constitute any prop-erty that we do ascribe to a token for semantic purposes from the other inferential properties of the token. A consideration of analogous demands elsewhere shows that this demand should be dismissed. A property may be constituted localistically out of some properties and not out of others. That may be the way the world is and nothing more needs to be said. (*b*) It is more likely that the demand for a principled basis is making a "basic" point. It raises the question: What makes a property that we ascribe for semantic purposes – a particular set of inferential properties – *a meaning?* We must distinguish the inferential properties of a token that are *meaning* constituters from the other inferential properties of the token. We do need a principled basis here. And we have one. A property – hence the inferential properties that constitute it – is a meaning if and only if it plays a semantic role and so is one we should ascribe for semantic purposes.

We are left with an epistemic problem: *showing* that localistic properties not holistic ones meet this criterion and are meanings. Three arguments are urged in the rest of the chapter. First, applying the "ultimate" method, all the properties we do ascribe for semantic purposes are in fact localistic. So, given the success of our current ascriptions in serving those purposes, we have good reason to suppose that the properties we ought to ascribe are localistic. Second, *in general,* whether our purposes are explana-tory, practical, or perhaps even frivolous, we tend to ascribe properties that are localistic because only localistic properties have the sort of generality we are interested in; localistic properties are

likely to be shared by many things. This yields the simplest, least
theory–laden, argument against semantic holism: We ought to
ascribe localistic properties because only such properties have the
generality that will serve our semantic purposes. Hence, only
localistic properties play semantic roles. Hence, all meanings are
localistic. Third, the popular, overarching theory, "Representa-
tionalism," that word meanings are entirely constituted by refer-
ential properties, provides a further argument, for no such mean-
ing is holistic.

In Chapter 4, I argue for a certain Representationalist program.
Applying the "ultimate" method, we find the descriptive version
of that doctrine confirmed by the classic discussion, generated by
Quine, of transparent and opaque ascriptions. The folk seem to
ascribe at least three different sorts of putative referential meaning
to a definite singular term: the property of referring to a specified
object under a specified mode (opaque ascription); the property
of referring to a specified object ("simply–transparent" ascription);
the property of referring en rapport to a specified object
("rapport–transparent" ascription). Given the success of ordinary
ascriptions in serving our semantic purposes, this is evidence that
we ought to ascribe these properties and, hence, that Representa-
tionalism is correct as a normative/basic doctrine too. Exploration
of the ways in which these ascriptions serve those purposes con-
firms this.

What we most need in order to explain these referential mean-
ings are theories of reference. I argue that three sorts of theory of
reference are possible. "Description" theories are one sort, but
they could not be true for all words. Some words must be covered
by "causal" theories – historical, reliablist, or teleological – ex-
plaining reference not in terms of the reference of other words
but in terms of direct noninferential relations to reality. And some
words may be covered by "descriptive-causal" theories.

I argue that the meanings (opaquely) ascribed to words may be
constituted by descriptive modes of reference: This is a familiar
molecularism. But I also argue for the apparently radical thesis that
some such meanings must be *causal* modes of reference. I illustrate

this with a historical-causal theory of names and other singular terms.

I reject rival programs. 'Fido'-Fido theories, hence direct-reference theories, of names fail because of well-known problems, particularly the identity problem: The true identities *'a = a'* and *'a = b'* differ in meaning. I do not argue for this difference in the usual way, by appealing to the differing informativeness of the identities, for that argument assumes Cartesian access to meanings. I argue for it by applying the methodology: We distinguish *'a'* from *'b'* in serving our semantic purposes, and we are right to do so. The evidence does not support "semi-Representationalist" two-factor theories as descriptive theories and counts against "anti-Representationalist" verificationist, use, and one-factor functional-role theories. The burden of showing that theories of these sorts are nevertheless normatively correct is very great.

I consider the meanings of attitude ascriptions – "second-level" meanings. I reject the view that these are extremely context dependent. I find support for an "intimate link," usually identity, between the meaning ascribed and a meaning of the ascribing content sentence. I argue that we should "put metaphysics first" in discussing ascriptions and hence take them to concern concrete thoughts and utterances rather than Platonic propositions. Finally, I develop the program to handle various puzzles including those due to Richard, Castañeda, and Kripke.

Chapter 5 is concerned with eliminativism and revisionism. Eliminativism is the view that nothing has a meaning. I take this to be an empirical doctrine that is not open to dismissal by popular transcendental arguments to the effect that the doctrine is "incoherent." Nevertheless, I think that its evidential support is weak.

Revisionism rejects the status quo: We ought to ascribe for semantic purposes properties other than the ones we do ascribe. Given the arguments in Chapter 4, I take the status quo to be Representationalist. I defend this position from two arguments, the argument from the computer analogy and the argument from methodological solipsism. Neither argument supports the view

that psychology should ascribe only syntactic properties, strictly understood, to mental states. The mind is not purely syntactic at any level. The argument from methodological solipsism may seem to support the view that psychology should ascribe only narrow meanings.

To assess this support we need to distinguish two views of narrow meaning. According to one, the narrow meaning of a sentence is a function taking an external context as argument to yield a wide meaning as value. According to the other, the narrow meaning is a functional role involving other sentences, proximal sensory inputs, and proximal behavioral outputs. Narrow meanings as functions must be acceptable to someone who believes in wide meaning. And they would yield explanations of behavior. I argue, however, that the moderately revisionist idea that we should ascribe these meanings instead of wide ones is mistaken; in particular, the meanings, and the behavior they would explain, are too coarse grained to serve our purposes. I am much more critical of the more popular functional-role narrow meanings. I argue that they are unexplained and mysterious. Even if they were not, we have been given no idea how such meanings *could* explain intentional behaviors. If they do not explain these behaviors, then revisionism requires that intentional behaviors be denied altogether. We have been given no reason to believe such a denial. These failings are very bad news for the highly revisionist doctrine that psychology should ascribe these putative meanings. That doctrine has a heavy onus arising from the apparently striking success of our present practice of ascribing wide meanings to explain behavior.

1

A critique of the
case for semantic holism

I. INTRODUCTION

1.1. Semantic Holism and Semantic Localism

At its most extreme, semantic, or meaning, holism is the doctrine that all of the inferential properties of a token in language or thought constitute its meaning. This doctrine is opposed by semantic localism, which, at its most extreme, denies that any of the inferential properties of a token constitute its meaning.

Despite its prima facie implausibility, semantic holism is ubiquitous. It has, as Jerry Fodor says, "something of the status of the received doctrine in the philosophy of language" (1987: 57). And it is urged, or taken for granted, in psychology and artificial intelligence. In this chapter, I shall look critically at the case for semantic holism.

The case can always be made to fit the following "basic" argument:

(1) Some of a token's inferential properties constitute its meaning.
(2) If some of a token's inferential properties constitute its meaning then they all do.
(3) So, all of a token's inferential properties constitute its meaning.

Fodor is an extreme "atomistic" localist: He resists this argument by rejecting premise (1) (pp. 73–95). Fodor's major reason for rejecting (1) is quite clear: He thinks that it leads inexorably to holism, which he regards as "a *crazy* doctrine" (p. 60) threatening Life As We Know It. He thinks that (1) has this unfortunate

10