

Michael Devitt is a distinguished philosopher of language. In this new book he takes up foundational issues in semantics.

Three important questions lie at the core of this book: What are the main objectives of semantics? Why are they worthwhile? How should we accomplish them? Devitt answers these "methodological" questions naturalistically and explores what semantic program arises from the answers. The approach is anti-Cartesian, rejecting the idea that linguistic or conceptual competence yields any privileged access to meanings.

This new methodology is used first against holism. Devitt argues for a truth-referential localism and in the process rejects direct-reference, two-factor, and verificationist theories. The book concludes by arguing against revisionism, eliminativism, and the idea that we should ascribe narrow meanings to explain behavior.

A substantial contribution to the literature on meaning and intentionality, this important study will be of particular interest to philosophers of language and mind and could be used in graduate-level seminars in these areas. The book will also appeal to readers in linguistics and the other cognitive sciences.

Published by the Press Syndicate of the University of Cambridge
The Pitt Building, Trumpington Street, Cambridge CB2 1RP
40 West 20th Street, New York, NY 10011-4211, USA
10 Stamford Road, Oakleigh, Melbourne 3166, Australia

© Cambridge University Press 1996

First published 1996

Printed in the United States of America

Library of Congress Cataloging-in-Publication Data

Devitt, Michael, 1938-

Coming to our senses : a naturalistic program for semantic
localism / Michael Devitt.

p. cm. - (Cambridge studies in philosophy)

Includes bibliographical references and index.

ISBN 0-521-49543-1. - ISBN 0-521-49887-2 (pbk.)

1. Semantics. 2. Semantics (Philosophy). 3. Holism. I. Title.
II. Series.
P325.D47 1996

121'.68-20

95-6275

CIP

A catalog record for this book is available from the British Library.

ISBN 0-521-49543-1 hardback

ISBN 0-521-49887-2 paperback

Contents

Preface	vii
Introduction	1
1. A Critique of the Case for Semantic Holism	10
I: Introduction	10
II: The Argument from Confirmation Holism	17
III: The Argument from the Rejection of Analyticity	18
IV: The Argument from Psychological Explanation	38
V: The Argument from Functionalism	43
2. The Methodology of Naturalistic Semantics	48
I: Introduction	48
II: Semantic Tasks	54
III: A Semantic Methodology	70
3. A Case for Semantic Localism	87
I: A Principled Basis	87
II: Definitions and Plan	95
III: Holism and Localism Elsewhere	102
IV: Holism and Localism in Semantics	113
4. Meanings and Their Ascription	136
I: A Representationalist Program	136
II: Rival Programs	171
III: Attitude Ascriptions	196
IV: Developing the Program	218

5. Eliminativism and Revisionism	245
I: Introduction	245
II: Syntactic Psychology	258
III: Narrow Psychology	285
References	313
Index	329

Preface

Two things led to this book. The more immediate, but less important, cause was my concern about semantic, or meaning, holism. Holism 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. Yet it seemed to me, as it did to Fodor, clearly false (“crazy” was his word). So, in 1989, I set out to show this.

First, I had to show that the arguments for holism were no good. The main argument stems from Quine: The localist idea that *some but not all* inferential properties of a token constitute its meaning (or content) is alleged to yield an analytic-synthetic distinction with epistemologically objectionable consequences. You can accept this argument without becoming a holist, of course, if you are prepared to adopt an “atomistic” localism according to which *no* inferential property *ever* constitutes the meaning of a token. That is Fodor’s path. However, atomism strikes me as implausibly extreme. Very likely, the meanings of some tokens are atomistic, but surely the meanings of others – perhaps ‘bachelor’ is an example – are not. I want to defend a “molecular” localism, according to which *a few* of the inferential properties of a token *may* constitute its meaning. I think that I can have what I want because I reject the Quinean argument: Molecular localism does *not* have epistemologically objectionable consequences unless it is saddled, gratuitously, with an epistemic thesis. I also reject other arguments against there being a “principled basis” for the molecular localist’s distinction among inferential properties.

It is one thing to reject arguments against there being a principled basis, it is another to show that there is one. This was the second thing I had to do to refute holism and establish localism. I found doing this much harder than I expected.

Attempting to do it soon raised some very general questions. What are the semantic tasks? Why are they worthwhile? How should we accomplish them? I have been bothered by these “methodological” questions from my semantic beginnings, long before holism marred my horizon. Signs of this bother are to be found scattered through my earlier writings.

That brings me to the less immediate, but more important, cause of this book: my desire to confront the methodological questions directly and thoroughly and to see what semantic program would follow from my answers. This cause is largely independent of the one arising from holism but, as I have indicated, not entirely so. My first use of my methodology is to show that localism has the principled basis it needs. So the program I urge is localistic.

Some parts of my semantic program – for example, truth-referentialism – are common enough. Some parts – for example, molecular localism – are not common at all (although I am here, unusually, in agreement with Michael Dummett). Some parts – for example, the view that a token has more than one meaning – are positively rare. I think that the main interest of this book lies not so much in these substantive theses as in its proposal of a novel methodology and in its use of this to support those theses.

I had a problem finding a title. A straightforwardly descriptive title using some combination of such well-worn terms as ‘meaning’, ‘mind’, ‘semantic’, ‘reference’, ‘representation’, ‘content’, ‘truth’, even ‘holism’ or ‘localism’, and certainly ‘realism’ or ‘reality’ would make the eyes glaze (and likely raise the question, “Wasn’t that one of Putnam’s books?”). I sought the help of my colleague, Michael Slote. In describing the book to him, I said that it proposes a methodology that it then uses against some popular but, in my view, very implausible views in semantics. Among these are holism, “two-factor” theories, and verifica-

tionism. Among them also is the “‘Fido’-Fido,” or “Millian,” theory of names, recently resurrected by “direct-reference” philosophers. According to this theory, a name lacks a sense, its meaning being simply its property of referring to its bearer. Finally, I emphasized to Slote that the semantic program proposed in the book was intended to be naturalistic and hence “derived from experience.” Putting all this together, in a divine moment of inspiration while “in the bathroom,” he thought of the nice pun “Coming to Our Senses.” (I should add that the senses I come to for names are not Fregean; they are nondescriptive senses made up of causal modes of reference.) The subtitle was easy and all my own work.

Work on the book led to two “unpublications,” “Meaning Localism” and “What Did Quine Show Us about Meaning Holism?,” and then to several publications. These are “Localism and Analyticity” (1993a), which is a brief early version of some ideas in Chapters 1 and 3; “A Critique of the Case for Semantic Holism” (1993b), of which Chapter 1 is a modified version; “Semantic Localism: Who Needs a Principled Basis?” (1994a), which is a brief early version of some ideas in Chapters 2 and 3; and “The Methodology of Naturalistic Semantics” (1994b), of which Chapter 2 is a modified version.

These papers, as well as draft chapters, have been delivered at many conferences and universities over the last five years. The book has benefitted from the ensuing discussions. It has also benefitted from the comments on some of these items of at least the following (my apologies to those omitted): Louise Antony, Ned Block, David Braun, Alex Byrne, Fiona Cowie, Rey Elugardo, Hartry Field, Jerry Fodor, Peter Godfrey-Smith, Dorothy Grover, Gil Harman, David Lewis, Eric Lormand, Peter Ludlow, Bill Lycan, Graham Oppy, Greg Ray, Michael Slote, Kim Sterelny, Corliss Swain, Ken Taylor, and, especially, Georges Rey.

I am grateful to the University of Maryland for giving me time off from teaching in the following ways: a Semester Research Award from the General Research Board of Graduate Studies and

Research in fall 1989; a Fellowship at the Research Center for Arts and Humanities in fall 1990; and a sabbatical leave in spring 1994. I am also grateful to the Australian National University for a Visiting Fellowship, March-June 1994.

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.

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 competence consists in knowledge of truth conditions. Another is the received Fregean view that two expressions that differ in informativeness 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 competence 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 empirical 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, “molecular,” 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 conventional wisdom that molecular localism is untenable because there is “no principled basis” for its distinction among inferential properties.

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 “Representationalist”: 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 Representationalism 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 revisionism 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 resurrected 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 meanings 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 semantics 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 property 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 explanatory, 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, “Representationalism,” that word meanings are entirely constituted by referential properties, provides a further argument, for no such meaning 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 Representationalism is correct as a normative/basic doctrine too. Exploration of the ways in which these ascriptions serve those purposes confirms this.

What we most need in order to explain these referential meanings 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, reliabilist, or teleological – explaining 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.

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

consequence because he accepts (2). Indeed, he is as committed to (2) as the most fervent holist.

I agree with Fodor's view of the holistic conclusion but think that he is quite wrong about (2). My aim in this chapter is to reject the case for (2). This aim is very important if we are to be safe from holism. First, the total rejection of (1) is difficult to sustain. Even if it is plausible that some tokens do not depend for their meanings on inferential properties, it is surely also plausible that some tokens do so depend; the likes of 'bachelor' spring to mind. If a significant proportion do so depend, and (2) is accepted, we still face a disagreeably holistic future. Second, we lack any persuasive argument against (1).¹ Finally, the localist should not put all her eggs in the one basket of rejecting (1). So I shall be defending a moderate "molecular" localism according to which a *few* of the inferential properties of a token *may* constitute its meaning.²

Why do people believe (2)? Premise (2) is accompanied by thinking along the following lines:

There is no principled basis for the molecular localist's distinction between the few inferential properties of a token alleged to constitute its meaning and all its other inferential properties. Only a token that shared all the inferential properties of the original token would really share a meaning with it.³

- 1 Fodor and Lepore (1992), in effect, offer an argument against *the case for* (1), but that is not of course the same as offering an argument against (1). (The qualification "in effect" is made necessary by their unusual presentation of the basic argument in terms of their notion *atomistic*.)
- 2 It is worth noting that a "cluster" localist (see sec. 3.4) might accept (1) without thinking that any one inferential property is *essential* to the meaning: What is essential is (a weighted) most of a small set of such properties.
- 3 If this sort of thinking were supported by a "slippery-slope" or "sorites" argument it could perhaps be swiftly dismissed. But it is not obvious that any of the arguments I shall consider are slippery slopes: It is not obvious that they rest on the claim that there is no distinction *because no sharp line can be drawn between the properties that count and the properties that do not* (cf. Fodor and Lepore 1992: 25). (Thanks to Ned Block and Georges Rey.)

So, to understand why people are semantic holists, we need to discover why this “no-principled-basis consideration,” offered in support of (2), seems plausible. I shall discuss and reject four arguments for the consideration: the argument from confirmation holism (part II), the argument from the rejection of analyticity (III), the argument from psychological explanation (IV), and the argument from functionalism (V).

Rejecting arguments that there is no principled basis does not of course establish that there *is* a principled basis. Some people are drawn to holism because they think “no one has provided a convincing reason for including some inferences and excluding others” (Block 1991: 40). I shall not be addressing this concern in this chapter; my arguments here are against the *arguments for* holism. I argue against holism itself in Chapter 3, where I claim that the demand for a principled basis is partly misconceived and that, insofar as a basis is needed, we have it.

Some preliminaries are necessary before considering the arguments.

1.2. Preliminaries

1. Two points about usage. (i) We need a convenient term to refer to the linguistic and the mental items that concern the dispute over semantic holism. I have chosen to use ‘token’ rather than ‘representation’ because it is less theory laden. (ii) The “meaning” of a token is the same as its “content.”

2. A thought is an attitude toward a mental token. I assume that the meaning of this token, like the meaning of a sentence, is complex. Furthermore, I assume that the mental token is also like a sentence in being itself complex: It has its meaning in virtue of its parts having meanings.⁴ Neither of these assumptions is

⁴ This is a further assumption because something without this complexity can have a complex meaning; for example, a yellow flag on a ship used to mean **THIS SHIP HAS YELLOW FEVER**. (I shall follow the convention of using uppercase bold words to refer to meanings and concepts.)

controversial in the holism dispute that concerns me. It would be a little controversial to go still further, however, and assume that the mental token is complex *in the way that a sentence is* (rather than, say, in the way a map is). This would assume the “language-of-thought” hypothesis. I shall argue for this assumption later (sec. 4.4), but it is no part of my argument against holism. Nevertheless, I shall apply the term ‘sentence’ to the full mental token in a thought, and ‘word’ to its meaningful parts, through lack of any better terminology.

3. A sentence could have a holistic meaning in virtue of having a holistic structure or holistic word meanings. I shall be criticizing the arguments for semantic holism about words.

4. I have located the difference between the holist and the localist in their views of the extent to which the *inferential properties* of a token contribute to its meaning. Sometimes the difference is located in the extent to which *beliefs* associated with the expression contribute to its meaning. Although these two versions of the dispute are in fact distinct, as we shall see (1.8), they are usually treated as equivalent. My purpose is to reject the case for both versions of holism. However, I shall mostly discuss “the inference version” and treat “the belief version” separately only when that is called for. The inference version of localism is close to the sort of semantics I shall defend.

5. The tokens we are concerned with are datable, placeable mental states, inscriptions, sounds, and so on. These tokens are, of course, *meaningful*, differing from meaningless ones (e.g., tokens written in the sand by the wind) in their historically given causal relations to speakers and hence to other meaningful tokens and the world. So they should not be thought of as if they were stripped of these relations and “uninterpreted.” By talking of meaningful tokens, we can avoid talk of “propositions.” This is an advantage not simply because propositions are creatures of darkness but also because talk of propositions in this context is explanatorily unhelpful (4.12). However, at one place I shall apply the discussion to propositions (1.10).

6. A mental sentence token has its inferential properties in

virtue of its actual and potential inferential relations to other sentences. These inferential relations are causal (and nonnormative) relations of a certain sort. The inferential properties of a mental word token are those that the sentence containing it has in virtue of that word (rather than in virtue of other words or solely in virtue of structure). The inferential properties of a linguistic token are derived from the properties of the mental token that caused it. It is likely that every word is related by inferential properties to every other word.

7. According to atomistic localism, a word's meaning is constituted not by its inferential properties but only by its links to the world (or to proximal sensory inputs and/or behavioral outputs).⁵ We shall see that such links must have *something* to do with the word's meaning, either directly, or indirectly via inferential relations to other words, in *any* plausible theory (4.5). So we should not take the holist as proposing that all meanings are *fully* constituted by inferential properties. The idea that the external relations of language have nothing to do with meaning is absurd (despite its popularity in the structuralist tradition).⁶

8. Talk of "the meaning" of a token is vague. I think that this vagueness plays an important role in the holist's misconception of the issue. In particular, the implication that a token has just one meaning is very misleading. A defense of localism will require our being much more precise in the next chapter, particularly about the purposes for which we attribute meanings. Nevertheless, in assessing the arguments for holism in this chapter, we can leave the talk of meaning vague.

5 The links will be to the world if the meaning is "wide," but they will be to proximal inputs and outputs if the meaning is "narrow," supervening on what is internal to the organism.

6 Fiona Cowie (1987) points out that holists are extremely casual about the place of extralinguistic links in constituting meaning, often writing as if the links had no place, as if meaning is constituted solely by the relations between tokens. The casualness she points to is part of a general tendency in the philosophy of psychology to ignore inputs and outputs (5.6–5.8).

9. Two of the four arguments for holism that I shall consider are due to Quine. However, I shall not consider his famous argument for the indeterminacy of translation (1960). Nor shall I consider arguments that start from a "principle of charity."⁷ I have argued against such principles elsewhere.⁸

1.3. A Straw Man?

The conclusion of the basic argument is that all of the inferential properties of a token constitute its meaning. This is a startlingly "individualistic" doctrine (in a non-Burgesian sense). For a token in my head to have the same meaning as one in yours there must be not the smallest difference in their inferential relations; in that respect, we must be functional duplicates. Indeed, for a token in my head this week to mean the same as a token in my head last week I must be a functional duplicate of my earlier self. As a result, it is almost certain that no person ever shares a single thought with any other person nor even with himself at a different time.

The individualistic aspect of this extreme holism may seem so startling as to raise doubts that anyone subscribes to it. Indeed, consider the recent response of Todd Jones, Edmond Mulaire, and Stephen Stich to Fodor's individualistic characterization of holism:

We can't think of anyone who has explicitly endorsed this very radical version of Holism. Nor is Fodor much help on this score; he

7 E.g., Davidson 1980: 239; 1984: 199–201; Harman 1973: 14; Putnam 1983: 149–50; 1988: 8–9. See Fodor and Lepore (1992: 59–104) for a criticism of such arguments. Sometimes the arguments seem not to be for semantic holism. They seem, rather, to presuppose semantic eliminativism (nihilism) – that there is (near enough) no fact of the matter about meaning – and to be for a holistic account of *our practice of ascribing meanings* (see, e.g., Putnam 1983: xiii). See section 2.7 for more on this.

8 1981a: 115–8; 1991b: 192–9; Devitt and Sterelny 1987: 244–9.

offers no references... Fodor's Meaning Holist is a straw man. (1991: 69)

Are they right?

A problem with semantic holism is that it is seldom stated clearly and explicitly. Fodor complains of this (1987: 55), and Jones, Mulaire, and Stich concede the point (1991: 69). Nevertheless, sometimes we find fairly explicit statements of an individualistic doctrine like extreme holism:

What our words mean depends on *everything* we believe, on *all* the assumptions we are making. (Harman 1973: 14)

every theoretical difference between individuals creates differences in the identities of their concepts and threatens reference failure whenever those theories are faulty. (Papineau 1987: 98)

the holistically individuated conceptual role of 'polio' in some world-view can not only be altered by finding out about the nature of the underlying etiology of the disease, but also by finding that, say, my Aunt Sally had it as a young girl. (McClamrock 1989: 260; the conceptual role is the "meaning" of 'polio')

if I say "Water is more greenish than bluish," and you say "Water is more bluish than greenish," then we have different narrow contents for "water." . . . in the real world we can expect no two cases to be subsumed by the same law of content. (Block 1991: 40-1)

Apart from what is explicit, something close to the extreme doctrine is implicit in most holist writings.⁹ One persuasive reason for thinking this is that if the arguments for holism were good then, with the exception of the argument from functionalism (V), they would establish an extreme doctrine.

Anything close to the extreme doctrine alarms Fodor and should alarm us all. However, my rejection of the arguments for holism does not depend on taking their conclusions as extreme. The arguments are not good in any case.

9 If not something more extreme: Many seem to take even differences of affective tone to be relevant to meaning.

II. THE ARGUMENT FROM CONFIRMATION HOLISM

1.4. The Argument

One cause of the inexplicitness of holism is the idea that Quine somehow established the doctrine years ago in "Two Dogmas" (1953: 20-46). What Quine established in "Two Dogmas" and elsewhere, to my satisfaction at least, was the *epistemological* view, confirmation holism: Put extremely, perhaps too extremely, the justification of a sentence depends on the justification of every other sentence. Combining this with the *semantic* view, verificationism, does indeed yield a simple argument for semantic holism. For, according to verificationism, the meaning of a sentence is its method of justification. So if that justification depends on every sentence then the meaning does as well.¹⁰

This combination of confirmation holism and verificationism can be related to the basic argument as follows. Consider the first premise of the argument:

(1) Some of a token's inferential properties constitute its meaning.

Why should we believe this? The verificationist has an easy answer: because it is (partly) in virtue of its inferential properties that a sentence is justified. If this is the reason for believing the first premise then confirmation holism leads to the second:

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

There can be no principled basis for distinguishing the inferential properties that count toward meaning from those that do not.

Despite its evident appeal, verificationism is not supported by

10 Quine 1960: 12-13; 1969: 80-1; 1981: 70-1. (Quine regrets the suggestion in "Two Dogmas" that the unit of significance is the whole of science: A substantial body of theory is sufficient; see 1991: 268.) There are signs of the argument in Putnam 1983: 144-7. Fodor summarizes the argument (1987: 62-3); so does Loar (1982: 273), attributing it to Harman.

any compelling argument and is, in my view, false (4.10). That it leads to semantic holism is a compelling reason to reject it (cf. Quine 1969: 81).¹¹

III. THE ARGUMENT FROM THE REJECTION OF ANALYTICITY

1.5. Analyticity and the Fregean Assumption

Independent of any verificationist assumptions, Quine is usually thought to have supported (2) and the no-principled-basis consideration in his attack on the analytic-synthetic distinction.¹² It is claimed that to reject (2) (after accepting (1)) is to support this distinction. As a result, some sentences would be known a priori; they could not conceivably be false; they would be unrevisable in that they must be held true come what may in experience; there would be privileged knowledge. Furthermore, the analytic-synthetic distinction would require principled distinctions between “change of belief and change of meaning,”¹³ between “collateral information and the determinants of content,”¹⁴ between “what properly belongs in a dictionary and what properly belongs in an encyclopedia.”¹⁵ Quine is supposed to have shown us that there is no such privileged knowledge and that there are no such distinctions.

An analytically true sentence is often said to be one that is true solely in virtue of meaning. A synthetically true sentence, in contrast, is one that is not true solely in virtue of meaning. What

11 Fodor and Lepore (1992: 37–58) argue that there is another crucial flaw in Quine’s argument: Confirmation holism *presupposes* semantic localism.

12 See, e.g., P. M. Churchland 1979: 46–54; P. S. Churchland 1986: 265–7; Putnam 1988: 8–11.

13 See Harman 1973: 108–9.

14 See Block 1986: 629.

15 Wilson 1967: 63. Wilson credits the formulation to Jerrold Katz. See also Harman 1973: 97–100; Block 1991.

has this distinction to do with the rejection of (2)? Suppose that a molecular localist rejects (2) while accepting (1) for some token of, say, ‘bachelor’. This has the consequence that the meaning of the token is constituted by its inferential links to tokens of some words¹⁶ – for example, ‘unmarried’ – but not to others – for example, ‘frustrated’. It may then look as though the sentence ‘All bachelors are unmarried’ is true solely in virtue of this fact about meaning and hence is analytic in the sense just noted.¹⁷ In contrast, ‘All bachelors are frustrated’ is not true solely in virtue of meaning and hence, if true at all, is synthetically so.

Appearances are deceptive. First, it has been insufficiently noted that *this line of reasoning depends crucially on the following “Fregean assumption”*: that inferential properties constitute meaning only insofar as they determine reference.¹⁸ From this assumption it follows that the meaning-determining link of ‘bachelor’ to ‘unmarried’ is also *reference* determining, and so ‘bachelor’ can refer only to objects that ‘unmarried’ refers to, that is, to unmarried objects. This is what makes ‘All bachelors are unmarried’ appear analytic. Without this assumption the meaning-determining link would be quite compatible with the sentence’s falsity.¹⁹

I agree with the Fregean assumption (Chaps. 4 and 5). As a

16 I shall mostly not be as pedantic as this about the type-token distinction.

17 What would appear to be analytic for a “cluster localist” (n. 2) would be something of the form ‘All *F*’s have (a weighted) most of the properties, *G*, *H*,’ And the same goes for the “cluster holist.” Cluster holism differs from cluster localism simply in the size of the set of properties in the cluster. I ignore cluster views in this chapter (but see Chap. 3). Taking account of them would require heavy qualifications in the discussion in this and the next section.

18 Note that inferential properties *could not* (fully) constitute the meanings, hence determine the references, of *all* words. Direct links to the world must (at least partly) do this job for some words (1.2, point 7; 4.5).

19 Putnam’s theory of “stereotypes” (1975: 139–52) provides a nice example of a semantics that rejects the Fregean assumption. Putnam thinks that an association with ‘yellow’ is part of the stereotype, and hence meaning, of ‘lemon’ even though ‘all lemons are yellow’ is not analytic or even true.

result, I think that a truth-referential semantics will explain the contribution of inferential properties to meaning at the same time that it explains reference. However, many disagree. They think that the inferential properties constitute one factor of meaning, and reference and truth constitute another. These two relatively independent factors require distinct semantics, a “narrow” functional-role (or “conceptual-” or “inferential-” role)²⁰ semantics *as well as* a “wide” truth-referential semantics.²¹ In this “two-factor” theory, the rejection of (2) does not imply any doctrine of analyticity, for that rejection is concerned with the inferential factor of meaning, whereas analyticity, being a doctrine of truth, is concerned with the other factor.

Paul Boghossian, however, has pointed out that we would miss the Quinean point if we concluded simply that the argument against analyticity has no holistic consequences for a non-truth-referential semantics.²² Any such semantics must explain meaning *some way*. So, it will be committed to some *analogue* of the Fregean assumption leading, apparently, to the view that some sentences have some *analogue* of analyticity. Thus, suppose that the semantics explains meaning not in terms of truth but in terms of warrant. Then it will be committed to the assumption that inferential properties constitute meaning only insofar as they determine warrant. So, for example, the meaning-determining link of ‘bachelor’ to ‘unmarried’ is also warrant determining. Then, the application of ‘bachelor’ is warranted only if the application of

20 I prefer “functional-role” for the reasons indicated later (4.9; 5.6–5.8).

21 See, e.g., Field 1977; Loar 1982; McGinn 1982; Block 1986. Earlier, Putnam had split the meaning of a word into one external component, the referent, and three internal components, one of which was a stereotype (1975: 269).

22 1993: 34. Boghossian is responding to Block’s claim, along the lines of the preceding paragraph, that narrow functional-role meanings are not analytic (1993: 18). My 1993b is mistaken in its handling of this matter (although, with Bill Lycan’s help, it went some of the way toward Boghossian’s point in n. 18).

‘unmarried’ is. It may then look as though ‘All bachelors are unmarried’ is warranted solely in virtue of meaning.

So it seems, after all, as if the molecular localist in rejecting (2) is committed via the Fregean assumption or an analogue to analyticity or an analogue. Two further points are worth making before moving to the second way in which appearances are deceptive.

I choose to defend a molecular localism committed to the Fregean assumption, but this is not an assumption that the holist should accept. For, consider its consequences for the reference of a word. According to the localist, if, through some error in a person’s theory, the conjunction of her word’s few reference-determining words fails to refer, so also does the original word; for example, if ‘adult unmarried male’ failed to refer, so would ‘bachelor’. According to the holist, *the slightest error in the person’s theory threatens general reference failure* (as David Papineau notes in the passage quoted in 1.3); the reference of each word depends on so many others in the theory. Given such truisms as “If ‘cat’ does not refer then there are no cats,” this loss of reference threatens *loss of the world*. Many holists resist this threat by claiming that we all “live in different worlds” of our theories’ making. Sensible holists will simply drop the assumption.

Whether we adopt the Fregean assumption or an analogue, it is odd to see the argument from the rejection of analyticity as an argument *for holism*. Suppose that the localist’s view that there is a meaning-determining relation between ‘F’ and ‘G’ commits her to the view that ‘All F’s are G’ is analytic (or an analogue). Then the holist should also be committed to this view. For the holist also thinks that there is that meaning-determining relation between ‘F’ and ‘G’. He differs from the localist in thinking that there are *many more* such relations: to ‘H’, ‘I’, ‘J’, and so on (some, presumably, yielding nonuniversal sentences). His disagreement with the localist should be simply over the *size* of the group of analytic (or analogue) statements. If commitment to analyticity (or analogue) posed a problem for the localist, it would pose a much worse one for the holist. I shall return to this.

1.6. Analyticity and Logical Truth

Appearances are deceptive in a second way because, even with the Fregean assumption, the sentence 'All bachelors are unmarried' is not true *solely* in virtue of meaning and so is not analytic in the sense discussed earlier. The sentence is indeed true partly in virtue of the fact that 'unmarried' must refer to anything that 'bachelor' refers to, but *it is also true partly in virtue of the truth of 'All unmarrieds are unmarried'*. The latter sentence is what Quine calls a "logical truth." So the molecular localist is committed to the sentence being analytic in another, perhaps less popular, sense: It can be "reduced by definition" to a logical truth. But her rejection of (2) does not commit her to the view that logical truths are true solely in virtue of meaning. So she is not committed to the sentence being analytic in the original sense.

In virtue of what is 'All unmarrieds are unmarried' true? The localist could, and I think should, answer as follows: It is true partly in virtue of what it means and partly in virtue of the way the world is, in virtue of all unmarrieds *being* unmarried. And because the truth of 'All bachelors are unmarried' depends on this logical truth, its truth *also* depends partly on the world.²³ So our localist can go along with the general Quinean dictum that the truth of every sentence depends partly on its meaning and partly on the world. Hence she does not believe that some sentences are true *solely* in virtue of meaning. She believes that they are true in virtue of meaning *given the logical truths*; they are analytic only in this weak sense.

'All bachelors are unmarried', logical truths, and mathematical truths are often thought to be *logically necessary*. The attempt to explain this necessity was a major motivation for doctrines of analyticity. It is important to note that the localist's rejection of (2) does not commit her to any doctrine with such pretensions.

23 I emphasize that this is not the uninteresting dependence of a sentence on the world for its *meaning*. With its meaning already fixed (1.2, point 5), the sentence depends on the world for its *truth*.

She might not accept that the truths are necessary. If she does think that they are, she need not accept any particular theory of that necessity.²⁴

1.7. The Argument from "Two Dogmas"

I have agreed that, in accepting (1) and rejecting (2) while accepting the Fregean assumption, the molecular localist is committed to a weak sort of analyticity. What has "Two Dogmas" got to do with this? The paper has two sorts of argument against analyticity. The first sort are arguments against attempts – mostly Rudolf Carnap's – to explain analyticity in terms of notions of *synonymy*, *state description*, *definition*, and *semantic rule*. Quine argues that the explanations do not break out of the intensional circle. But *these* arguments cannot, of course, prejudge *all* attempts to break out of the circle. They must leave it as an open empirical question whether a scientifically respectable account of meaning can be given, and whether a theory that gives this will be a molecular localist one with weak analyticity of the sort just noted. So I shall say no more of these arguments.

The second sort of argument has been more influential. It establishes certain *epistemological* views, in particular, confirmation holism and empirical revisability. A consequence is that no belief, not even a law of logic or mathematics, is immune to revision in the face of experience; the web of belief is "seamless." The contrasting view – that the web is seamed – supplied the other major motivation for traditional doctrines of analyticity. According to those doctrines, analytic beliefs have a privileged epistemic status; they are known a priori and are empirically unrevisable. So Quine's "Two Dogmas" stands clearly opposed to the epistemo-

24 Cf. Lepore and Fodor who argue (against my 1993a) as if I *must* use analyticity to explain my modal talk (1993: 674). If this point had been sound, it would have been much more effective against them: Their own modal talk would have committed them to analyticity, something they dread almost as much as holism.

logical aspects of traditional doctrines of analyticity. *But localism's rejection of (2) does not commit it to any of these epistemological aspects.*

The rejection of (2) is simply a *semantic* matter, making no epistemic claims at all. Insofar as the considerations just cited provide a route from this rejection to a doctrine worthy of the name "analyticity," it is a route to a nonepistemic doctrine.²⁵ This

25 Of course the doctrine *is about* inferential properties, which are clearly epistemic: Inferential properties are significant in belief formation. But the doctrine *is not itself* epistemic: It does not entail that the difference between inferential properties that constitute a meaning and ones that do not has epistemic significance; it does not entail anything about the epistemic status of sentences – for example, that some are known a priori.

This distinction between *being about* something epistemic and *being* something epistemic is important in assessing Fodor and Lepore's rejection of molecularism. Their official position is that Quine undermined doctrines with "an epistemic criterion like apriority or unrevisability" for analyticity (1992: 58), doctrines that *are* epistemic in the respect just noted. Despite this, they frequently argue as if Quine has *thereby* undermined doctrines that are merely *about* something epistemic, hence undermined *any* sort of molecularism (e.g., p. 57). This conclusion is in direct contrast to what I am arguing here. I made this argument against Fodor and Lepore briefly before (1993a). In their reply, they persist:

On our view, Quine brought into question the possibility of *any* reconstruction of 'meaning constituting inferences' in terms of 'inferences accepted'. The crucial point is that there is, according to Quine, no way of distinguishing the inferences X accepts *because of what he believes* from those he accepts *because of what he means*; . . . (Lepore and Fodor 1993: 673; see also Fodor and Lepore 1993b: 309)

But "the crucial point" is aimed at a doctrine that *is* epistemic – analyticity implying something like apriority about *the conditions under which* inferences are accepted – whereas the possibility allegedly "brought into question" is a doctrine that might be merely *about* something epistemic – a doctrine about meanings explained in terms of inferences accepted. So their point is not relevant, let alone crucial. The molecular localist who accepts (1) but rejects (2) is committed to nothing whatever about *why* inferences are accepted. Why do Fodor and Lepore think otherwise? The answer is clear: They are *totally* convinced that a nonepistemic criterion for analyticity cannot be found (see, e.g., the passage quoted in sec. 3.1). But a conviction is not an argument. They are entitled to demand that the molecular localist

doctrine will yield a privileged epistemic status to the sentence 'All bachelors are unmarried' *only if we add assumptions about our knowledge of meanings and about our knowledge of logical truths*. There is no compelling reason for the localist to make any assumptions contrary to Quine's naturalized epistemology.

Consider logical truths. The localist may think that these, like any other, are true partly in virtue of their meaning and partly in virtue of the way the world is. The logical truths do of course enjoy a privileged epistemic status of some sort, but there need be no more to this than Quine indicated: A sentence earns its place on the list of logical truths by having a certain centrality in our web of belief.

Consider meanings. Cartesianism is still rife in this realm. Linguistic (conceptual) competence is thought to give "privileged access" to meanings. Merely understanding the words 'bachelor' and 'unmarried' (merely having the concepts **BACHELOR** and **UNMARRIED**) yields propositional knowledge of their meanings, including the relations these meanings have to each other. And this knowledge can be brought before the conscious mind by "analysis." The localist need not go along with this Cartesianism. She can see competence not as semantic propositional knowledge but as an *ability or skill*: It is knowledge-how not knowledge-that. Not only can she have this view of competence, I shall argue that she should have this view (Chap. 2).

Once (1) is accepted, we have seen that the rejection of (2), which is so important to stopping holism, commits the Fregean molecular localist to the doctrine that some sentences can be "reduced by definition" to logical truths. This doctrine does not pretend to explain either logical necessity or apriority.²⁶ Perhaps, therefore, it is not worthy of the name "analyticity." That is of no

produce a nonepistemic criterion. I shall attempt to meet this demand in Chapter 3. But, if my present argument is correct, they are not entitled to give the impression that Quine has *already shown* that there can be no such criterion.

26 This weak doctrine is in the spirit of Putnam (1975: 33–40). See also Antony 1987.

concern to the localist. Her concern is simply to reject (2), not to resurrect a traditional doctrine of analyticity. I shall now draw out in more detail how this rejection is perfectly compatible with a Quinean epistemology. I shall consider the matter of a priori knowledge first, the matter of revisability second, and the matter of the unprincipled distinctions third.

1.8. A Priori Knowledge

A priori knowledge is often claimed to be *knowledge not derived from experience*, knowledge that is, or could be, justified without any appeal to experience. The Fregean molecular localist is not committed to such knowledge.

Suppose that Joe believes tokens of

(A) All *F*'s are *G*

(S) All *F*'s are *H*

– they are in his “belief box.” Suppose further that the meaning of ‘*F*’ depends on its inferential links to ‘*G*’ but not on its links to ‘*H*’. On the strength of this, let us say that (A) is, but (S) is not, “weakly analytic.” Suppose that Joe not only believes these sentences but *knows* them. *Must* the localist think that there is a difference in the way Joe knows, or could know, them? No. She may think that the processes by which they get into the belief box and are maintained there must both meet the same standard of justification, whatever that may be. The story for both sentences can be fully in accord with confirmation holism.

Of course, there is a route to knowledge of (A) that is not available for (S). If Joe knows (i) that the reference of ‘*F*’ in his belief box is (partly) determined by the reference of ‘*G*’, and (ii) that ‘All *G*'s are *G*’ is true, then he can infer (A). However, this route need not be *interestingly* different from the route for the nonanalytic (S). Both (i), which is an application of theoretical semantic knowledge, and (ii), which is a logical truth, may be arrived at in the usual empirical way.

Consider (i) in particular. The Cartesian view is that, as a result

of his linguistic competence, Joe (tacitly) *knows that* ‘*G*’ stands in the meaning-determining relation to ‘*F*.’ By taking competence as an ability or skill, the localist can reject this. So, competence with ‘*F*’ consists (partly) in being disposed to infer tokens of ‘*x* is *G*’ from tokens of ‘*x* is *F*’. It does not consist in knowing that ‘*F*’ is related in this way to ‘*G*’ nor in knowing that this relation (partly) constitutes the meaning of ‘*F*’. That knowledge would be an application of a semantic theory. Joe may well be in a privileged position to apply this theory to his own thoughts, for he has ready access to those thoughts. But this does not give the theory or its application any special epistemic status (2.11).

However, suppose I were wrong about this. Suppose that we became convinced by the analyticity argument that molecular localism has epistemologically unacceptable consequences. Then we should also be convinced that holism has even worse ones: If localism yielded privileged knowledge of a few facts, holism would yield privileged knowledge of many.

The Fregean molecular localist's rejection of (2) gives no support to the idea that there can be knowledge that is not derived from experience. But there is another idea of a priori knowledge (or, perhaps, another way of understanding the first idea): It is knowledge of the world that *can be gained in the process of learning a word or concept*. Let us explore this idea.

Earlier (1.2, point 4), I distinguished two versions of the holism-localism dispute. According to the inference version the dispute is over the extent to which the *inferential properties* of a word constitute its meaning. According to the belief version the dispute is over the extent to which *beliefs* associated with the word constitute its meaning. So far, we have been discussing the inference version although the arguments could be adapted to the belief version. At this point, it is necessary to attend briefly to the difference between the two versions.

Clearly, if it were the case that a belief in (A) constituted the meaning of the word ‘*F*’, then it *would follow* that Joe could not gain the word without gaining the belief. In contrast, if it were the case that the inferential practice of inferring ‘*x* is *G*’ from ‘*x* is

F' constituted the meaning, then it *would not follow*. To get from his inferential practice to the belief, Joe would need to go through a further process; for example, an inferential process involving the following three other beliefs: the belief that he followed the practice; the belief that the practice is good; and the logical belief that it would not be good unless (A) were true. This demonstrates that the two versions of the holism-localism dispute are not equivalent.

On the inference version, competence alone does not even require a *belief* in (A) and so there can be no question of it requiring knowledge of (A). On the belief version, competence does require the belief; but how could competence supply the justification that turns the belief into *knowledge*? Only, once again, by assuming a Cartesianism access to meanings and by taking knowledge of logical truths for granted. We have already seen that the localist need not accept that competence yields propositional knowledge about '*F*'s. No more must it yield propositional knowledge about '*F*'s.²⁷

1.9. Unrevisable Sentences

Here is a typical statement of a common worry:

analyticity has obvious epistemological consequences: An analytic sentence would be *unrevisable*, in the sense that to deny or reject it would be *eo ipso* to abandon its standard meaning; one who called it false would be, as Quine says, not denying the doctrine but changing the subject. Thus nothing could count as evidence against the truth expressed by an analytic sentence, and more generally we could have no rational grounds for doubting that truth (we could be mistaken only about the meanings of the relevant words). (Lycan 1991: 112)

I think that the idea that analyticity leads to an unacceptable unrevisability is misguided.

We should first remind ourselves that the localist is not committed to the *logical necessity* of analytic sentences (1.6). So if there

27 This goes against Devitt and Sterelny 1987: 79–80.

is to be a relevant worry about revisability it must concern *our epistemic relations* to these sentences.

Consider the inference version of localism. For Joe to revise his opinion of (A) is for him to drop it from his belief box. Our discussion shows that Joe *can* drop (A) without changing the meaning of '*F*'. For, belief in (A) is *not constitutive* of understanding the meaning of '*F*'. Thus, Joe might drop the sentence because he is unaware that he follows the practice of inferring '*x* is *G*' from '*x* is *F*', or because he does not realize that this practice is meaning constituting, or because he does not realize that there is a logical link between the practice and the sentence.

Perhaps it is not *psychologically* possible for Joe to drop (A) from his belief box without changing meaning. But nothing interesting follows from this, and certainly nothing un-Quinean does. It is, for example, quite compatible with Quinean epistemology that it should be psychologically impossible for humans to abandon various logical truths.

Suppose that Joe *knew* that (A) was weakly analytic. Then he would know that he can drop it from his belief box only by changing its meaning or abandoning a logical truth. But, of course, his knowledge that (A) is weakly analytic is as revisable as can be: It is the result of the fallible application of a fallible theory; there need be no question of it being held come what may in experience.²⁸

Revisability does not even *appear* to be a problem for the inference version of localism. However, it may appear to be a problem for the belief version. For, in that version, weakly analytic sentences are indeed not revisable without meaning change. I shall argue that this is not a problem. But suppose it were. It

28 Morton White pointed out years ago that "the statement 'All men are rational animals' is analytic' is itself empirical" (1950: 320). This point has not had the impact it deserves. Note that even if there *were* sentences that were *strongly* analytic in that they were true *solely* in virtue of meaning, hence not in virtue of anything about the extralinguistic world, they would still be revisable. For, one's opinion of the linguistic facts upon which such truths would depend would be revisable.

would be a worse problem for holism. For, where the localist makes a few beliefs unrevisable without meaning change, the holist makes many. *Worries about revisability provide no argument for holism.*

Here are some perfectly general considerations. A token is of a certain type – whether it be a cat, a pain, a hammer, a philosopher, a pawn, a capitalist, or whatever – in virtue of having certain properties. Those properties constitute its being of that type so that if something did not have those properties it would not be of that type. Often the constitutive properties are relational; they are ones an object has in virtue of its relations to other things.

Let us compare two tokens that are of relational types: a token that is a capitalist, and a token that means **BACHELOR**. The token capitalist has many relational properties that may change over time. Some of these changes, for example, ceasing to own a Volvo, do not affect the person's still being a capitalist. Others, for example, ceasing to own means of production, do affect this: If she loses that property, she ceases to be a capitalist. But there is nothing more in principle to stop her losing what is essential to being a capitalist from what is inessential. Similarly, a particular token in Joe's belief box that has the relational property of meaning **BACHELOR** may change over time. (Or, if the idea of such a token continuing through time is farfetched, think of it being replaced by another with different relations. I shall ignore this subtlety.) According to the localist, some of these changes do not affect the meaning of the token but others do. In the inference version, we have been supposing that ceasing to be inferentially related to tokens meaning **FRUSTRATED** does not affect meaning whereas ceasing to be inferentially related to tokens meaning **UNMARRIED** does. In the belief version, ceasing to be appropriately related to a token sentence in the belief box meaning **ALL BACHELORS ARE FRUSTRATED** does not affect meaning whereas ceasing to be appropriately related to one that means **ALL BACHELORS ARE UNMARRIED** does. But there is nothing more in principle to stop the token from losing what is essential to its meaning than from losing what is

inessential. Something cannot cease to be appropriately related to means of production and still be a capitalist. Something cannot cease to be appropriately related to certain other tokens and still mean **BACHELOR**.

In general, the nature of a type constrains which tokens can have it. When the type is that of being a capitalist, the constraints are economic. When the type is that of meaning **BACHELOR**, the constraints are semantic. The latter constraints should be no more shocking or surprising than the former, even when they involve beliefs.

According to a semantic theory, dropping a token that means **ALL BACHELORS ARE UNMARRIED** changes meanings. According to an economic theory, ceasing to own means of production changes something from being a capitalist. The former fact is no more reason for abandoning the semantic theory than is the latter for abandoning the economic theory. *Such unrevisability is not epistemological; it is harmlessly "metaphysical."*

There may, of course, be good semantic reasons for doubting (as I do) that the meaning of a token is partly constituted by a certain belief. My point is simply that the worry about unrevisability should not be among those reasons.

My argument involves a distinction between the properties of a token that constitute its being of a certain type and the properties that do not. This may seem to result in an un-Quinean involvement in modalities. But it need not. The distinction and the modalities can be sustained by mere regularities; for example, the fact that all capitalists own means of production but they do not all own Volvos. This is acceptable even to Quine (1966: 50–1). Apart from that, this modal issue is quite general, having no special bearing on semantics or epistemology.

Neither version of localism conflicts in any way with confirmation holism, even in the extreme form in which I stated it. In the face of experience, each token in the belief box, with its meaning-determining relations, might have to be assessed against any other one, with its meaning-determining relations. As a result of this assessment, any token could be dropped from the belief