A Critique of the Case for Semantic Holism

Michael Devitt

A CRITIQUE OF THE CASE FOR SEMANTIC HOLISM

Michael Devitt
University of Maryland

I. Introduction
I.1. Semantic Holism and Localism

At its most extreme, semantic, or meaning, holism is the doctrine that all of the inferential properties of an expression constitute its meaning. This doctrine is opposed by semantic localism which, at its most extreme, denies that any of the inferential relations of an expression 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" (Fodor 1987: 57). And it is urged, or taken for granted, in psychology and AI. In this paper 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 an expression's inferential properties constitute its meaning.
(2) If some of an expression's inferential properties constitute its meaning then they all do.
(3) So, all of the inferential properties of an expression constitute its meaning.

Fodor is an extreme localist: he resists this argument by rejecting premise (1) (pp. 73-95). Fodor's major reason for rejecting (1) is very 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 paper 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 hard to sustain. Even if it is plausible that some expressions do not depend for their meanings on inferential properties, it is surely also plausible that some expressions 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 localism according to which a few of the inferential properties of an expression may constitute its meaning.[2]

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

> The localist's distinction between the inferential properties that constitute an expression's meaning and the ones that do not has no principled basis. All of these properties bear on the semantic role of the expression. Perhaps some are more important than others, but importance is a matter of degree. There is no real distinction between the ones that count toward the meaning and the ones that do not. Only an expression that shared all the inferential properties of the original word would really share a meaning with it.[3]

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 consider 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 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 here; my arguments are against the *arguments for* holism not against holism. Elsewhere (forthcoming), I argue that the demand for a principled basis is largely misconceived and that, insofar as a basis is needed, we have it. I shall summarize this view at the end.

Some preliminaries are necessary before considering the arguments.

I.2. Preliminaries

1. My talk of "expressions" covers both the linguistic and the mental, the complex and the simple. I shall apply the terms 'sentence' and 'word' to mental

expressions—thoughts and their parts—as well as to linguistic ones. The "meaning" of an expression is the same as its "content."

2. I shall focus on semantic holism about words, for I take that to be more basic than holism about sentences.

3. I have located the difference between the holist and the localist in their views of the extent to which the *inferential properties* of an expression 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 (III.4), they are usually treated as equivalent. My purpose is to reject the case for both versions of holism. However, I shall mostly discuss "the inferential version" treating "the belief version" separately only when that is called for. I think that the inferential version of localism is close to the sort of semantics we should defend.

4. The talk of expressions in the basic argument is of *tokens*; of datable, placeable mental states, inscriptions, sounds, etc. These tokens are, of course, *meaningful*, differing from meaningless ones (for example, 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 them in this context is explanatorily unhelpful. However, at one place I shall apply the discussion to propositions (III.6).

5. A mental sentence token has its inferential properties in virtue of its actual and potential inferential relations to other sentences. These inferential relations are (nonnormative) causal 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.

6. The extreme localist believes that a word's meaning is constituted not by its inferential properties but only by its links to the world. Such links must have *something* to do with the word's meaning, either directly, or indirectly via inferential relations to other words, on *any* plausible theory.[4] 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).[5]

7. 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 would require our being much more precise, particularly about the

purposes for which we attribute meanings. Nevertheless, in assessing the arguments for holism, we can leave the talk of meaning vague.

8. 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."[6] I have argued against such principles elsewhere.[7]

1.3 A Straw Man?

The conclusion of the basic argument is that all of the inferential properties of an expression token constitute its meaning. This is a startlingly individualistic doctrine. For a token in my head to have the same meaning as one in yours there must not be the smallest difference in their inferential relations; we must be functional duplicates (so far as inference is concerned). 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 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). 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)
> ...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.[8] 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.

II. The Argument from Confirmation Holism

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 so does the meaning.[9]

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 an expression'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 an expression'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 any compelling argument and is, in my view, certainly false. That it leads to semantic holism is a compelling reason to reject it (cf Quine 1969: 81).[10]

III. The Argument from the Rejection of Analyticity
III.1 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.[11] 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,"[12] between "collateral information and the determinants of content,"[13] between "what properly belongs in a dictionary and what properly belongs in an encyclopedia."[14] 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 has this distinction to do with the rejection of (2)? The rejection of (2) has the consequence that the meaning of a token, say a token of 'bachelor,'[15] is constituted by its inferential links to some words—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 above sense.[16] 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.*[17] From this assumption it follows that 'bachelor''s meaning-determining link to 'unmarried' is also *reference*-determining, and so 'bachelor' can only refer to objects that 'unmarried' refers to; that is, to unmarried objects. This is what makes 'All bachelors are unmarried' appear analytic. Without the assumption the meaning-determining link would be quite compatible with the sentence's falsity.[18]

I agree with the Fregean assumption. As a 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 (1981, 1989). However, many disagree. They think that the inferential properties constitute one factor of meaning while reference and truth constitute another. These two relatively independent factors require distinct semantics, a "narrow" conceptual-role (or functional-role) semantics *as well as* a "wide" truth-referential semantics.[19] On 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. So, for this theory, the argument against analyticity could have no direct holist consequences. Yet, interestingly enough, two-factor theorists tend to be holists. Presumably they are influenced by other arguments.

The dependence of the argument from the rejection of analyticity on the Fregean assumption undermines it as an argument *for holism*. For, clearly, an

easy localist response to the argument is simply to drop the assumption. There are two reasons why the holist should not object to this response.

1. Suppose that the Fregean 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. Then the Fregean 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 in thinking that there are *many more* such relations: to 'H,' 'I,' 'J,' etc. (some, presumably, yielding non-universal sentences.) His disagreement with the localist should be simply over the *size* of the group of analytic statements. If commitment to analyticity posed a problem for the localist who accepts the Fregean assumption, it would pose a much worse one for the holist who accepts the assumption.

2. The holist should not accept the Fregean assumption in any case. 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 her theory threatens general reference failure* (as David Papineau notes in a passage quoted in I.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.[20]

According to received opinion, we must accept the no-principled-basis consideration and holism to avoid analyticity. But localism could imply analyticity only given the Fregean assumption. So localists could avoid analyticity by dropping the assumption, as many holists do and all should. *The argument from the rejection of analyticity is really an argument against the Fregean assumption not localism.* However, the assumption is worth defending and so we must proceed.

There is another reason for proceeding: one aspect of the analyticity argument—the revisability worry—can be resurrected without the Fregean assumption. Yet, we shall see that even if that aspect were good *the argument would still do nothing to show that non-Fregean holism is preferable to non-Fregean localism* (III.5).

III.2 Analyticity and Logical Truth

Appearances are deceptive in a second way. 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 above sense. 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 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 a re 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, the fact that all unmarrieds *are* unmarried. And since the truth of 'All bachelors are unmarried' depends on this logical truth, its truth *also* depends partly on the world.[21] 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. 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.

III.3 The Argument from "Two Dogmas"

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 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 moderately localist one with weak analyticity of the above sort. 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 epistemological aspects of traditional doctrines of

analyticity. *But localism's rejection of (2) does not commit it to any of these epistemological aspects.*[22]

The rejection of (2) is simply a *semantic* matter, making no epistemic claims at all.[23] Insofar as the considerations above provide a route from this rejection to a doctrine worthy of the name "analyticity," it is a route to a nonepistemic doctrine. This 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 good reason for the localist to make any assumptions that yield unQuinean conclusions.

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 make meanings transparent to the mind. 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. Furthermore, the Cartesianism is implausible and almost entirely unargued, apparently the result of the beguiling use of the vague term 'know' in the ordinary locution, 'knows the meaning.' I have argued this at length elsewhere.[24]

Once (1) is accepted, we have seen that the rejection of (2), which is so important to stopping holism, commits the Fregean 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 aprioricity.[25] Perhaps, therefore, it is not worthy of the name "analyticity." That is of no 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.

III.4 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 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 he knows, or could know, them? No. 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 is not *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, must 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 *knows that* 'G' stands in the meaning-determining relation to 'F.' The localist should reject this, seeing competence not as semantic propositional knowledge but as a skill: it is knowledge-how not knowledge-that. So competence with 'F' consists (partly) in tending 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.' This 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.

    I have pointed out that the analyticity argument is really against the Fregean assumption not localism (III.1). The point is worth emphasizing here. On the one hand, if the assumption is rejected, as it is by two-factor theories, then the concern about a priori knowledge is *obviously* irrelevant to localism. For, without the assumption, the inferential links between words need not yield truths. For example, even if mere understanding of the words 'lemon' and 'yellow' *did* yield propositional knowledge of their meanings leading, perhaps, to the belief, 'All lemons are yellow,' that would not establish that the belief was true.[26] Without the Fregean assumption, knowledge of meanings could not yield knowledge of facts and so could not yield *privileged* knowledge of the facts. On the other hand, if the assumption is accepted and we become convinced by the analyticity argument that localism has unacceptable epistemological con-sequences, 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 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 above 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 (I.2), I distinguished two versions of the holism-localism dispute. According to the inferential 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 inferential version although the arguments could be adapted to the belief version.[27] 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) were constitutive of 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' were constitutive, it *would not follow*. To get from his inferential practice to the belief, Joe would need two other beliefs: the belief that he followed this practice and the logical belief that the practice would not be good unless (A) were true. This demonstrates that the two versions of the holism-localism dispute are not equivalent.

On the inferential 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 there is no way that competence alone can supply the justification that turns the belief into *knowledge*. We have already seen that competence does not yield propositional knowledge about 'F''s. No more does it yield propositional knowledge about F's.[28]

Perhaps this is too hasty. Georges Rey has suggested to me that, where (A) is indeed weakly analytic, a belief in it gained in the process of acquiring 'F' might count as one known a priori. For, it might be claimed, this way of gaining it is *reliable*. But, as Rey points out, if this is a priori knowledge it is harmless. For, Joe might know something a priori in this sense and yet not know that he does. To know that the knowledge is a priori, Joe has to know and apply a semantic theory; he has to know that the sentence is weakly analytic. And that knowledge is, of course, as empirical as any.[29] So, the fact that a sentence is known a priori in this sense does not supply any basis for foot-stamping dogmatism about it. Necessarily, if it is known a priori, then it is true. But, necessarily, if it is known at all, then it is true. The key empirical question remains: Is it known, whether a priori or not? In brief, this sort of a priori knowledge supplies no privileged knowledge.

III.5. 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 revisability is
misguided.

We should first remind ourselves that the localist is not committed to the
logical necessity of analytic sentences (III.2). So if there is to be a relevant
worry about revisability it must concern *our epistemic relations* to these
sentences.

Consider the inferential 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 meaning-determining practice
of inferring 'x is G' from 'x is F'; or because he does not realize that there is a
logical link between this 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 unQuinean 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 only drop it from his belief box by changing its meaning or abandon-
ing 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 is no question of it being held come what may in experience.[30]

Revisability does not even *appear* to be a problem for the inferential version
of localism. However, it may appear to be for the belief version. For, on 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. There
would be two interesting consequences.

First, we should note that the problem would remain even if we dropped the
Fregean assumption. Without that assumption the localist would not, of course,
be committed to any analyticity. However, on the belief version, she would still
be committed to some sentences being unrevisable without meaning change. If a
belief in (A) is constitutive of the meaning of 'F' then Joe cannot drop (A) from

his belief box without changing meaning whether (A) is analytic or not. So, to this extent, we can resurrect the analyticity argument without the Fregean assumption. Second, with or without that assumption, if revisability were a problem for the localist it would be a worse problem for the holist. 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 a little far-fetched, 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. On the inferential version, we have been supposing that ceasing to be inferentially related to tokens meaning FRUSTRATED does not affect meaning whilst ceasing to be inferentially related to tokens meaning UNMARRIED does. On 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 whilst 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 losing what is essential to its meaning than what is inessential. Something cannot both cease to be appropriately related to means of production and be a capitalist. Something cannot both 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 unQuinean 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 we stated it. In the face of experience, each token in the belief box, with its meaning-determining relations, might have to be assessed against every other one, with its meaning-determining relations. As a result of this assessment, any token could be dropped from the belief box. Linguistic or conceptual competence alone brings no knowledge (or, following Rey's suggestion, no knowledge of knowledge) and so does not prevent this revision.

Localism can accept that a person could be wrong about anything: she can believe false sentence tokens and disbelieve true ones.[31] If it seems inconceivable that we could be wrong about some sentence, that needs explanation. Perhaps the weak analyticity of the sentence could be part of the explanation. That would require a psychological link between that analyticity and the inconceivability (a link that we would have to discover without making the assumption that competence alone yields propositional knowledge of meanings.) And there are other possible explanations of inconceivability as Quine, Putnam, and others have shown. In any case, the inconceivability of error cannot establish truth; what was once inconceivable is often now taken for granted, as many familiar examples show.

In sum the inferential version of localism does not even appear to have a problem with revisability. The belief version may appear to have but does not really: the unrevisability it is committed to has no consequences that are epistemologically or otherwise objectionable.

III.6 Unrevisable Propositions

So far I have avoided any talk of propositions. Some will think that this misses the main point. They will object that it is a consequence of localism that

some *propositions* are unrevisable and that this is dreadful. But they would be wrong. Our discussion carries straight over to propositions.

On the basis of the inferential links between 'F' and 'G,' and the logical truth of the proposition that all F's are F, we might say that the proposition that all F's are G is weakly analytic. Is this proposition unrevisable? We should start by reminding ourselves, once again, that the localist is not committed to any logical necessities. So the worry about revisability must concern our epistemic relations to analytic propositions.

For Joe who believes the proposition that all F's are G to cease to believe it is for him to drop from his belief box any sentence that expresses the proposition. Obviously Joe can do this. But suppose he *entertains* that proposition, then surely he must believe it? Not on the inferential version of localism: entertaining a sentence that expresses that proposition requires engaging in certain inferential practices but it does not require any particular belief. On the belief version, in contrast, entertaining the proposition does indeed seem to require believing it. But Joe is no more prevented from revising his relation to the proposition than was our capitalist from revising her relation to the means of production. If Joe is to entertain the proposition he must believe it. If the person is to be a capitalist she must own means of production. In neither case is this "metaphysical" unrevisability objectionable.

Could Joe come to believe the proposition that some F's are not G? Certainly he could on the inferential version. On the belief version, the answer is not so clear. For Joe to have a token that means SOME F'S ARE NOT G in his belief box he would also have to have one that means ALL F'S ARE G. Perhaps the localist will answer that this is not possible. Or perhaps she will answer that in those circumstances Joe would believe neither the proposition that some F's are not G nor the proposition that all F's are G. More plausibly, if less charitably, she may answer that Joe would believe both propositions. (Uncertainty about the answer explains the cautious "seem" in the last paragraph.) Whatever the answer, there is no cause for localist concern.

The unrevisability worry about localism is a storm in a teacup. And if it were not it would be more of a worry about holism.

III.7. The Unprincipled Distinctions

Quine is thought to have shown not only that analyticity has undesirable epistemic consequences, but also that it involves unprincipled distinctions. Once again, Quine's argument does not undermine localism.

(a) The inferential version of localism involves no unprincipled distinction between "collateral information and the determinants of content" because it is not committed to *any* information being determinant of content. (b) Similarly, this localism does not involve any unprincipled distinction between "change of belief and change of meaning" because change of meaning is not a matter of changing

belief but a matter of changing inferential practices. (c) Finally, this localism does not involve a *disturbing* distinction between encyclopedia entries and dictionary definitions. If the dictionary definitions are taken to use rather than mention words, then in *some* cases,[32] they may differ from encyclopedia entries in being weakly analytic. However, such definitions will still be informative.

The belief version of localism does involve a distinction between collateral information and the determinants of content. And it has the consequence that some changes of belief change meaning and some do not. But Quine's argument does not count against these consequences. The concern about a priori knowledge is irrelevant because the distinctions do not lead to a priori knowledge. And the revisability worry no more shows that these distinctions are unprincipled than it shows that the distinction between the properties that constitute being a capitalist and the properties that do not is unprincipled.

III.8. Conclusion

The argument from the rejection of analyticity is not an argument against localism and for the no-principled-basis consideration and holism. Strictly, it is an argument against the Fregean assumption that what constitutes meaning determines reference. In any case the analyticity argument fails. The Fregean localist is only committed to weak analyticity. Unless she also adopts some epistemological theses that are both unnecessary and mistaken, this analyticity has no disturbing epistemic consequences. It involves no commitment to knowledge not derived from experience. There may be a sense in which knowledge can be obtained in gaining a concept, but it is a harmless sense. There is no objectionable respect in which knowledge is unrevisable, to be held onto come what may in experience, even on the belief version of localism (which I do not recommend). The rejection of (2) is a semantic matter, not an epistemological one.

IV. The Argument from Psychological Explanation

The belief version of the first premise in our basic argument for holism, is:

(1)* Some of an expression's associated beliefs constitute its meaning.

Why believe this? It seems to be common to believe it because of the role of expressions in psychological explanation. We must posit a difference in the meaning of two expressions in order to explain differences in behavior that result from them. It is thought that this difference in meaning can only be accounted for by a difference in associated beliefs. Some are tempted to go further:

(2)* If some of an expression's associated beliefs constitute its meaning then they all do.

The holist conclusion follows.

I think that many arguments in the psychological literature come perilously close to this. Consider also the following:

> Individuals with different theoretical assumptions are led thereby to accept different beliefs and to perform different actions. Why then insist that such individuals have the same causal-role concepts, when the tendencies to thought and action we want those concepts to inform us about are different?" (Papineau 1987: 98).

Finally, Ned Block's adaptation (1991: 60-1) of Putnam's Ruritanian example (1983: 144-7) is an argument of this sort.

I shall bring out the flaws in this sort of argument by considering Block's version in some detail. The major flaw is the final move to (2)*. This simply *assumes* what the holist is supposed to be arguing for, the no-principled-basis consideration. However, I think that this unargued move may seem plausible because of the reasoning for (1)*: if that reasoning worked for some beliefs it may seem obvious that it would work for all. So I shall focus on showing that that reasoning is also flawed.

There are two stages to Block's argument:[33]

> **Stage one:** Barry and Bruce are 10-year-old twins who were adopted at birth into different homes in North and South Ruritania, where they each observe adults drinking a liquid called 'grug' that causes the drinkers to act silly. Northern "grug" is Scotch whisky, southern "grug" is beer. Though there are many differences between Scotch and beer, none of these differences has *as a matter of fact* differentially impinged on Barry and Bruce. All that Barry and Bruce know about their respective "grugs" is that they are roughly earth toned liquids that make the adults drunk. Now clearly their uses of 'grug' differ in reference and hence in "wide" meaning. However, since there are no relevant differences between Barry and Bruce "in the head," their uses of 'grug' are alike in "narrow" meaning. And narrow meaning is what we must advert to for the purposes of psychological explanation.
> **Stage two:** Bruce and Barry reach adolescence and learn four ordinary but different facts about what they call 'grug.' For example, Bruce learns,

> (B)   It takes a lot of grug to make you drunk.

and Barry learns,

> (W)   It doesn't take much grug to make you drunk.

> Now the behaviors of Bruce and Barry will clearly differ as a result of these changes. So the needs of psychological explanation require that we see their uses of 'grug' as now differing in narrow meaning. Yet the only changes from stage one are the acquisitions of these beliefs. So we must see these beliefs as constituting the respective meanings of 'grug.' But

there is nothing special about any of these beliefs: many other equally run-of-the-mill beliefs would have done as well. Every such change in belief must be seen as changing the meaning of 'grug.' Indeed, the narrow meaning of 'grug' continually changes with the acquisition of knowledge.

(i) I shall not challenge the assumption that narrow meaning is what we need for psychology. And I am not concerned to argue against (1)* —indeed, I think that the closely related (1) is probably true for some expressions—nor to argue against the claim that the uses of 'grug' by Bruce and Barry at stage two have different narrow meanings. However, it is important to see that the above considerations fall far short of establishing these latter two claims. It is indubitable that the expected difference between the behavior of Bruce toward beer and Barry toward whisky must be explained by a difference in the meanings of their expressions. But the story itself supplies an obvious explanation: the *non-* 'grug' difference between, for example, (B) and (W). A person who believes that a lot of x but not much y makes you drunk will behave differently toward x and y *whether his concepts of x and y are the same or different*. Where there is a difference in *conception* we have no need to hypothesize a difference in *concept* in order to explain behavioral difference.[34] We cannot establish, for example, that my 'tiger'-concept is different from my 'koala'-concept *simply* by pointing out that I run away when I think that a tiger is on the loose but not when I think that a koala is (cf. Block 1991: 61). For that difference is sufficiently explained by the fact that I think tigers but not koalas are dangerous. To assume that a difference in conception entails a difference in concept is to assume holism not argue for it.

(ii) But suppose that after the acquisition of a few beliefs like (B) and (W) Bruce and Barry really do differ in their narrow 'grug'-meanings. *Must* we say that this is attributable to those beliefs, thus accepting (1)*? Of course, if those beliefs really are the only changes that have taken place, then we must. But other changes are possible. I made the point earlier that direct links to the world are important to meaning (I.2, 6). The point arises from the needs of reference determination and so is about wide meaning, but a similar point can be made about narrow meaning. First, suppose we accept the Fregean assumption. Then a word's narrow meaning is a function that yields a wide meaning as a value given the external context as an argument.[35] So, the narrow meanings of some words (at least) must be partly (at least) constituted by the fact that, given certain contexts, their wide meanings are partly (at least) constituted by certain direct links to that context. Next, suppose we reject the Fregean assumption, as Block does. Then, the narrow meanings of some words (at least) must be functional roles partly (at least) constituted by direct links to stimuli and/or behavior.[36] Either way the changes in the meanings of 'grug' may be attributable to changes in these extra-linguistic links—changes resulting perhaps from experiences of

beer and whisky —not to the changes in beliefs. This issue cannot be settled by consulting our intuitions about the meagre facts provided in examples like this. Settling it requires the application of a theory that accommodates the vast range of evidence relevant to a theory of meaning.

(iii) But suppose, finally, that some of these associated beliefs do indeed constitute the respective meanings. Where is the argument that they all must? Where is the argument for (2)*? On the Fregean assumption, for example, the beliefs that constitute meaning determine reference and so are in that respect "special." It does not follow that all beliefs determine reference and are special. Of course, it may be felt that there is no principled basis for distinguishing the ones that determine from the ones that do not. But, to repeat, *that is to assume holism not to argue for it.*

In sum, psychological considerations of the sort adduced by Block do not show that the adolescent Bruce and Barry differ in their narrow 'grug'-meanings. Even if they do so differ, that difference may not be attributable to differences in belief. Even if it is so attributable, no reason has been offered for thinking that all changes in belief change meanings. The argument does not establish holism; it begs the question.

V. The Argument from Functionalism

Acceptance of (1) amounts to acceptance of a conceptual-role, or functional-role, element to meaning. According to the Fregean assumption, this element determines reference. So it should be accommodated within a truth-referential semantics. This is the view I favor. However, many do not, claiming instead that a special functional-role semantics is required (III.1). The received view then seems to be that this semantics is essentially holistic. So we must accept (2). Thus Fodor claims that "functional-role semantics...is...inherently holistic" (1987: 83). Stich assumes without argument that the functional-roles that constitute his "fat syntactic" properties are holistic (1991: 248-9). Ron McClamrock simply *identifies* an individualistic doctrine of the most alarmingly extreme sort (I.3) with functional-role semantics (1989: 260).

Why do people think that functional-role semantics must be holistic? They may of course be influenced by the arguments that we have considered so far. However, there is often no sign of this. Furthermore, the most popular argument, that from the rejection of analyticity, certainly *should not* be influential because it depends on the Fregean assumption which functional-role semantics rejects. I think that a further argument is influential: *functional-role semantics is functionalist and functionalism is essentially holistic.* Functionalism is thought to give no principled basis for distinguishing the functional roles that are constitutive from the ones that are not.

How holistic is functionalism? I shall answer this using functionalist theories of the mind as examples.

According to David Lewis (1983: 99-107), what makes something the type *pain* is that its tokens typically have the causes and effects of pains captured by the platitudes of folk psychology; and then what makes something a token pain is that it is a token of that type. So this functionalist theory is, in some sense, holistic. I shall say that it is "type"-holistic. We shall see that type holism is very mild.

We do not rest with folk psychology: we go scientific. As a result, we hope to come up with fresh generalizations. Do we then *have to* make these constitutive too and hence *have to* replace our old concept PAIN with a new one PAIN*? *Must* we take our new theory not to be telling us a lot that we didn't know about pains but rather to be telling us about something different, but perhaps related, pains*? In brief, is functionalism *essentially* type-holistic? The received answer to these questions seems to be, "Yes." Thus, Block, in a helpful account of functionalism, feels obliged to talk not of pain but of pain "relative to theory T"; every theory must have its own concept of "pain" (1980b: 174). When Fodor says, "psychofunctionalism type-individuates mental states by reference to the psychological generalizations that subsume them" (1987: 70), he means *all* the generalizations. Bill Lycan seems to think that *all* the relations of a mental state must go into characterizing it (1988: 49).

It is implicit in these discussions of type holism that it is only the *noticed* relations of mental states that constitute their natures: platitudes and psychological generalizations are things we theorists believe. So any relations that pains or pains* may have that we have not yet discovered are not constitutive of their natures as pains or pains*.

It is implicit also in these discussions that the relations that constitute mental states are ones that are realized in a population including at least all (normal) people: the platitudes and generalizations cover us all. So Nigel, Lawrence, and Anna can all have pains or pains* even though what Nigel has makes him think of Eton, what Lawrence has turns him on, and what Anna has makes her angry. *There is nothing individualistic* about the type holism of functionalist mental properties (nor of functionalist economic properties, to take another example; see the discussion of *being a capitalist*, III.5 above).

The extreme semantic holism that most alarms us differs strikingly from type holism in being individualistic: it is not a holism of *shared* properties; all the inferential relations of a token, whether shared or not, constitute its meaning (I.3). So, although the differences between Nigel, Lawrence, and Anna do not prevent them sharing pains, they do prevent them sharing meanings. This is the most implausible and the most damaging feature of extreme holism. It gets no support from functionalism.

The individualism comes from taking the basic argument for holism to concern expression *tokens*. I think there is good reason to think that this is what

most holists intend (I.3). However, suppose we take the argument to concern types. Then its holistic conclusion is that all the inferential properties of an expression type—all the properties typical of tokens of that type—constitute its meaning. This is a move toward the type holism of Lewis' theory of the mind but it still differs from that holism in an important aspect: it requires that all the *actual* inferential properties of a word constitute an expression's meaning, whether we have noticed them in our theory or not. This aspect gets no support from functionalism.

Finally, if premise (2) of the basic argument—now taken to concern types—is to be supported by functionalism *alone*, then it needs to be established that functionalism is *essentially* type-holistic. If it is simply an accidental matter of fact that some functionalist theories are holistic, we need further argument to show that functional-role semantics is one of those holist theories. The view that functionalism is essentially holistic seems to be common, but it is implausible and, so far as I know, unargued.

Suppose, nevertheless, that functionalism were essentially holistic. Then functional-role semantics would be *type*-holistic. The mildness of this "social" holism becomes apparent when we notice that only a few of a token's inferential properties *are* likely to be believed typical of tokens of its type. Thus, it may well be the case that the only inferential properties believed typical of tokens that mean BACHELOR are ones that relate those tokens to ones that mean ADULT, mean UNMARRIED, or mean MALE. Such properties are likely to be only a few of the inferential properties of any given token that means BACHELOR. So type holism may be compatible with moderate localism.

In sum, functionalism gives no support to extreme semantic holism. And it gives no support to any holism at all without an argument to show that it is essentially holistic. Even if this were shown, the resulting semantic holism would be very mild.

Some holists seem to find the no-principled-basis consideration *obvious*, too obvious to really need an argument. Functionalism should give these holists pause. If the consideration were obvious, why would it not be equally obvious that there was no principled basis for saying that Nigel, Lawrence, and Anna all have pains?

VI. Conclusion

Semantic holism rests on (2) and the no-principled-basis consideration. I have examined four arguments for holism and found them all wanting. The argument from confirmation holism fails because it rests on verificationism. The argument from the rejection of analyticity is not really for holism but rather against the Fregean assumption. It fails anyway because it saddles the localist with unacceptable epistemic assumptions. The argument from psychological ex-

planation fails because it begs the question. Finally, the argument from functionalismcould notestablish an extreme holism and, in any case, is incomplete.[37]

To complete the case against semantic holism, we need to present the positive case for localism. I am attempting to present this case elsewhere (forthcoming). I summarize.

The appeal of semantic holism arises from the failure to situate the issue properly. Meanings are properties we theorists and folk attribute to tokens for theoretical and practical purposes. The variety of our purposes leads us to ascribe more than one such semantic property to a token. In these two respects, semantic properties are like the properties we ascribe elsewhere: in biology, economics, and so on. Semantic holism's idea that tokens *can have* only holistic semantic properties is misconceived. It has no more to be said for it than the analogous idea that a token *can have* only holistic biological properties. And we need no "principled basis" to show this. Whatever constitutes the natures of properties, whether in semantics, biology, or wherever, can constitute both holistic and localistic ones. The interesting issue is not this ontological issue but one about theories: *Should we ascribe* holistic or localistic semantic properties to tokens? Evidence that we should ascribe localistic ones is provided by the fact that the properties we *do* ascribe in ordinary 'that'-clauses are localistic. Further evidence is provided by the fact that the practice of ascribing localistic properties is completely general, to be found in biology, economics, ordinary life, or wherever. Finally, this general practice is as "principled" as could be: *localistic properties serve our theoretical and practical purposes, in particular our interest in generality.* We could ascribe holistic properties but it would be pointless to do so. An interesting question remains: Why does ascribing one localistic property *rather than another* serve our purposes? In semantics, I think that the answer is to be found by looking to what determines reference.[38]

Notes

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 *anatomistic*.)
2. It is worth noting that a "cluster" localist 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 amounted to 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.)
4. Must the meaning involve links to *the world* rather than merely links to *our experiences of* the world? It must because otherwise we should have to either leave reference out of meaning or overlook the message of Putnam's slogan, "meanings just ain't in the head" (1975: 227). I defend the slogan in my 1990.

5. Fiona Cowie (unpublished) points out that holists are extremely casual about the place of extra-linguistic links in constituting meaning, often writing as if they had no place: meaning is constituted solely by the relations between expressions. The casualness she points to is part of a general tendency in the philosophy of psychology to ignore inputs and outputs (Devitt 1989; 1991a).

6. 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, for example, Putnam 1983: xiii).

7. 1981: 115-8; 1991b: 192-9; Devitt and Sterelny 1987: 244-9.

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

9. 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.

10. Fodor and Lepore (1992: 37-58) argue that there is another crucial flaw in Quine's argument: confirmation holism *presupposes* semantic localism.

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

12. See Harman 1973: 108-9.

13. See Block 1986: 629.

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

15. I shall mostly not be as pedantic as this about the type/token distinction.

16. What would appear to be analytic for a "cluster localist" (note 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 what follows. Taking account of them would require heavy qualifications in points 1 and 2 below.

17. Where inferential properties constitute the meaning of a word, they make that meaning dependent on the meanings of other words. So those properties could not (fully) constitute the meanings of *all* words. Direct links to the world must (at least partly) constitute the meanings and determine the references of some words (I.2, point 6; Devitt and Sterelny 1987: 51-2).

18. The appearance of analyticity depends on the Fregean assumption if truth is explained in terms of reference. But what if truth is explained in terms of warranted assertability? (Thanks to Bill Lycan.) The appearance then depends on an analogous assumption: inferential properties constitute meaning only in so far as they determine warranted assertability. So, strictly speaking, either the Fregean assumption or some analogue is necessary for the appearance of analyticity.

19. See, for example, Field 1977; McGinn 1982; Loar 1982; Block 1986. Earlier, Putnam had split the meaning of a word into three internal components, including a "stereotype," and one external component, the referent (1975: 269). The stereotype does not determine reference.

20. Could we argue similarly that the holist should not accept the analogue of the Fregean assumption set out in note 18? It seems not, but the identification of truth with warranted assertability threatens loss of the independent world anyway (Devitt 1991b: 44-6).

21. I emphasize that this is not the uninteresting dependence of a sentence on the world for its *meaning*. With its meaning already fixed (I.2, point 4), the sentence depends on the world for its *truth*.

22. Fodor and Lepore emphasize that Quine's argument counts only against an analytic/synthetic distinction with "an epistemic criterion" like aprioricity (1992: 56-8). However, they presuppose that no other criterion is possible and hence argue that Quine has refuted moderate localism (p. 57).

23. Of course (2) is *about* inferential properties which are clearly epistemic: inferential properties are significant in belief formation. The point is that neither (2), nor its rejection, are epistemic. The distinction between inferential properties that constitute meaning and ones that do not—a distinction that comes with the rejection of (2)—need have no epistemic significance at all: it entails nothing about the epistemic status of any sentence.

24. 1981: 95-110; 1983: 674-5; 1991b: 270-5; Devitt and Sterelny 1987: 146-8; 1989.

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

26. The example comes from Putnam's original discussion of stereotypes (1975: 139-52): he thinks that the 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.

27. For example, we could follow the lines of the argument in III.1 to show that the apparent analyticity of 'All bachelors are unmarried' depends on an adapted Fregean assumption.

28. This goes against Devitt and Sterelny 1987: 79-80.

29. 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.

30. 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 extra-linguistic world, they would still be revisable. For, one's opinion of the linguistic facts upon which such truths would depend would be revisable.

31. On the belief version, of course, a token of hers could not *both* be disbelieved by her *and* have a meaning that is weakly analytic (perhaps, for example, the meaning that all bachelors are unmarried). But note the need for caution in the next section.

32. But certainly not all: many dictionary entries are obviously not candidates for analyticity.

33. This is a modified and shortened version. I mean to capture the essence of his argument.

34. My conception of echidnas is a set of beliefs involving the concept ECHIDNA. So is yours. Unless holism has been established, we have no reason to reject the familiar view that these very different conceptions involve the same concept. Block himself notes the concept/conception distinction earlier (1991: 49) but ignores it in discussing Ruritania.

35. I think that this is all that needs to be said about narrow meaning, thus agreeing with Fodor 1991 against Block 1991 (although I am being more tolerant than Fodor of the idea that associations with beliefs may constitute wide meaning and hence narrow meaning).

36. I have argued (1989) against Stich (1983) that the syntactic properties of a word are insufficient for psychological explanation.

37. I have not considered Stich's well-known discussion of Mrs T (1983: 54-8, 84-6). Yet Fodor claims that "the received view is that Mrs T makes a case for the holism of belief content" (1987: 62). If this is the received view it ought not to be. I have noted that Stich assumes without argument that his "fat syntactic" properties are holistic. These are properties that he thinks *cognitive science ought to* (perhaps does) ascribe. His discussion of Mrs T is part of one concerned with the very different matter of what *the folk do* ascribe (1983: 73-110). When a person ascribes a belief, Stich argues, she ascribes something similar in certain respects to a certain belief of her own, the required similarity varying holistically

with the context. He does seem to conflate this holism with the holism of content (pp. 54, 106), but *there is no clear argument for the latter*; i.e. for the view that the content that is appropriately similar in the context is constituted by a large proportion of the belief's associations. Holistic ascription of content is one thing, holistic content ascribed is another.

38. Most of the main ideas in this paper were in the first draft of "Meaning Localism," an unpublished 1989 paper. They were also in "What did Quine Show us about Meaning Holism?", delivered at the annual conference of the Society for Philosophy and Psychology in San Francisco in June, 1991. I am indebted to the following for comments: Louise Antony, Ned Block, Alex Byrne, Fiona Cowie, Hartry Field, Jerry Fodor, Gil Harman, David Lewis, Bill Lycan, Corliss Swain, and especially Georges Rey.

References

Antony, Louise M. 1987. "Naturalized Epistemology and the Study of Language." In *Naturalistic Epistemology*, edited by Abner Shimony and Debra Nails. Dordrecht: D. Reidel Publishing Company: 235-57.

Block, Ned, ed. 1980a. *Readings in Philosophy of Psychology, Volume 1*. Cambridge, Mass.: Harvard University Press.

Block, Ned. 1980b. "Introduction: What is Functionalism?" In Block 1980a: 171-84.

Block, Ned. 1986. "Advertisement for a Semantics for Psychology". In French, Uehling and Wettstein 1986: 615-78.

Block, Ned. 1991. "What Narrow Content is Not." In Loewer and Rey 1991: 33-64.

Churchland, Patricia S. 1986. *Neurophilosophy: Toward a Unified Science of the Mind-Brain*. Cambridge, Mass.: MIT Press.

Churchland, Paul M. 1979. *Scientific Realism and the Plasticity of Mind*. Cambridge: Cambridge University Press.

Cowie, Fiona. Unpublished. "Meaning Holism." BA thesis, University of Sydney, 1987.

Davidson, Donald. 1980. *Essays on Actions and Events*. Oxford: Clarendon Press.

Davidson, Donald. 1984. *Inquiries into Truth and Interpretation* Oxford: Clarendon Press.

Devitt, Michael. 1981. *Designation*. New York: Columbia University Press.

Devitt, Michael. 1983. "Realism and Semantics," part II of a critical study of French, Uehling and Wettstein 1980. *Nous* 17: 669-81.

Devitt, Michael. 1989. "A Narrow Representational Theory of the Mind." In *Rerepresentation: Readings in the Philosophy of Psychological Representation*, edited by Stuart Silvers. Dordrecht: Kluwer Academic Publishers: 369-402. Reprinted in *Mind and Cognition: A Reader*, edited by William G. Lycan. Oxford: Basil Blackwell.

Devitt, Michael. 1990. "Meanings Just Ain't in the Head." In *Meaning and Method: Essays in Honor of Hilary Putnam*, edited by George Boolos. Cambridge: Cambridge University Press: 79-104.

Devitt, Michael. 1991a. "Why Fodor Can't Have it Both Ways". In Loewer and Rey 1991.

Devitt, Michael. 1991b. *Realism and Truth*. Oxford: Basil Blackwell, 2nd edn revised (1st edn, 1984).

Devitt, Michael. Forthcoming. *Coming to Our Senses: A Program for Semantic Localism*.

Devitt, Michael, and Kim Sterelny. 1987. *Language and Reality: An Introduction to the Philosophy of Language*. Oxford: Basil Blackwell.

Devitt, Michael, and Kim Sterelny. 1989. "Linguistics: What's Wrong with 'the Right View.'" *Philosophical Perspectives, Volume 3: Philosophy of Mind and Action Theory, 1989*, edited by James E. Tomberlin. Atascadero, CA: Ridgeview Publishing Company.

Field, Hartry. 1977. "Logic, Meaning, and Conceptual Role". *Journal of Philosophy* 74: 379-409.
Fodor, Jerry A. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind.* Cambridge, Mass.: MIT Press.
Fodor, Jerry A. 1991. "Replies." In Loewer and Rey 1991: 255-319.
Fodor, Jerry A, and Ernest Lepore. 1992. *Holism: A Shopper's Guide.* Oxford: Basil Blackwell.
French, Peter A., Theodore E. Uehling Jr., and Howard K. Wettstein, eds. 1980. *Midwest Studies in Philosophy, Volume V: Studies in Epistemology.* Minneapolis: University of Minnesota Press.
French, Peter A., Theodore E. Uehling Jr., and Howard K. Wettstein, eds. 1986. *Midwest Studies in Philosophy, Volume X: Studies in the Philosophy of Mind.* Minneapolis: University of Minnesota Press.
Harman, Gilbert. 1973. *Thought.* Princeton: Princeton University Press.
Jones, Todd, Edmond Mulaire, and Stephen Stich. 1991. "Staving Off Catastrophe: a Critical Notice of Jerry Fodor's *Psychosemantics.*" *Mind and Language*, 6: 58-82.
Lewis, David K. 1983. *Philosophical Papers, Volume 1.* New York: Oxford University Press.
Loar, Brian. 1982. "Conceptual Role and Truth-Conditions". *Notre Dame Journal of Formal Logic* 23: 272-83.
Loewer, Barry, and Georges Rey, eds. 1991. *Meaning in Mind: Fodor and his Critics.* Oxford: Basil Blackwell
Lycan, William G. 1988. *Judgement and Justification.* Cambridge: Cambridge University Press.
Lycan, William G. 1991. "Definition in a Quinean World." In *Definitions and Definability: Philosophical Perspectives*, edited by J. H. Fetzer, D. Shatz and G. Schlesinger. Dordrecht: Kluwer Academic Publishers: 111-131.
McClamrock, Ron. 1989. "Holism Without Tears: Local and Global Effects in Cognitive Processes." *Philosophy of Science* 56: 258-74.
McGinn, Colin. 1982. "The Structure of Content". In Woodfield 1982: 207-58.
Papineau, David. 1987. *Reality and Representation.* Oxford: Basil Blackwell.
Putnam, Hilary. 1975. *Mind, Language and Reality: Philosophical Papers, Volume 2.* Cambridge: Cambridge University Press.
Putnam, Hilary. 1983. *Realism and Reason: Philosophical Papers, Volume 3.* Cambridge: Cambridge University Press.
Putnam, Hilary. 1988. *Representation and Reality.* Cambridge, Mass.: MIT Press.
Quine, W. V. 1953. *From a Logical Point of View.* Cambridge, Mass.: Harvard University Press.
Quine, W. V. 1960. *Word and Object.* Cambridge, Mass.: MIT Press.
Quine, W. V. 1966. *Ways of Paradox and Other Essays.* New York: Random House.
Quine, W. V. 1969. *Ontological Relativity and Other Essays.* New York: Columbia University Press.
Quine, W. V. 1981. *Theories and Things.* Cambridge, MA: Harvard University Press.
Quine, W. V. 1991. "Two Dogmas Revisited." *Canadian Journal of Philosophy* 21: 265-74.
Stich, Stephen P. 1983. *From Folk Psychology to Cognitive Science: The Case Against Belief.* Cambridge, Mass.: MIT Press.
Stich, Stephen P. 1991. "Narrow Content Meets Fat Syntax." In Loewer and Rey 1991.
White, Morton G. 1950. *John Dewey: Philosopher of Science and Freedom.* New York: The Dial Press.
Wilson, N. L. 1967. "Linguistic Butter and Philosophical Parsnips." *Journal of Philosophy* 64: 55-67.
Woodfield, A., ed. 1982. *Thought and Object.* Oxford: Clarendon Press.