

Peacocke, W. V. O. Quine, Georges Rey, Mark Rollins, Michael Root, Gideon Rosen, David Rosenthal, Stephen Schiffer, John Searle, Gabe Segal, Robert Stalnaker, and Steven Stich. We are likewise indebted to colleagues and students at Rutgers University, Washington University, and The City University of New York, where much of this material was developed in graduate seminars.

We wish to express our gratitude to Stephan Chambers and Andrew McNeillie for having seen us through, Luca Bonatti for his help with the references, Catherine McKeen for her help with the index, and to Jean van Altena for copy-editing. Special thanks to Dr Jessie Rosenthal.

A version of some material from chapter 6 was published in the 1991 *Proceedings of the Chicago Linguistics Society* under the title "Why meaning (probably) isn't conceptual role."

J.F.
E.L.

1

INTRODUCTION:
A GEOGRAPHY OF THE ISSUES

ANATOMIC AND HOLISTIC PROPERTIES

This is a book about semantic holism. Semantic holism is a doctrine about *the metaphysically necessary conditions* for something to have meaning or content. We therefore commence our discussion by attempting to view semantic holism in its metaphysical context.

Many properties have the property of being, as we shall say, *anatomic*.¹ A property is anatomic just in case if anything has it, then at least one other thing does. Consider, for an untendentious example, the property of being a sibling. If I am a sibling, then there is someone whose sibling I am; someone other than me, since no one can be his own sibling.² My being a sibling is thus, as one says, *metaphysically dependent* upon someone else's being a sibling (and so too, of course, is my sibling's being a sibling). So the property of *being a sibling* is anatomic; I *couldn't be* the only person in the world who instantiates this property. If I could *prove* that I am a sibling, that would refute solipsism.

If a property is not anatomic, then we shall say that it is *atomistic or punctate*. An atomistic or punctate property is one which might, in principle, be instantiated by only one thing. (So, for example, all properties expressed by predicates like "discovered the only . . ." or "ate the last . . ." are punctate, and so

too, we suppose, is the property of being a rock.) One way of formulating a main issue to be discussed in this book is whether *being a symbol, being a symbol belonging to language L, having an intentional object, having intentional content, expressing a proposition, having a referent, being semantically evaluable*, and the like are punctate properties. The currently received philosophical view is that these sorts of properties are not punctate but anatomic. We propose to explore the arguments for this view.

Many anatomic properties have the property of being *very* anatomic, or, as we shall say, the property of being *holistic*. Holistic properties are properties such that, if anything has them, then *lots* of other things must have them too. The “lots of” part of this definition could bear to be sharpened, no doubt; but, for our purposes, this isn’t required. Our primary concerns in this book will be with natural languages and with minds. Natural languages and minds can be assumed to be productive in all the interesting cases; minds (in any event, human minds) can grasp endlessly many different ideas, and languages (in any event, human languages) are capable of expressing endlessly many distinct propositions. The semantic properties we’ll discuss will therefore generally be ones which, if they are holistic, then if anything at all has them, so too do endlessly many other things.

Consider, for an untendentious example of a holistic property, *being a natural number*. Some philosophers have brought themselves to doubt that anything has this property; to doubt, that is, that numbers exist. For all we know, it is coherent – even well advised – of them to doubt this. But nobody could coherently doubt – and, so far as we know, nobody has ever sought to do so – that if there are any numbers, then there must be quite a few. One couldn’t, for example, coherently wonder whether there is only the number three.

Why not? Well, according to standard treatments, the natural numbers are defined by reference to the *successor relation*: nothing is a natural number unless there is a natural number

that is its successor. No number is its own successor, so if anything is a natural number, something else must be a natural number too; the existence of each natural number is thus metaphysically dependent on the existence of other natural numbers. That is, the property of being a natural number is anatomic. So far the number case is quite like the sibling case, but now the examples diverge. For whereas every sibling is his sibling’s sibling, no number is its successor’s successor (or its successor’s successor’s successor, and so forth). So, if there are any siblings, then there must be at least two of them; but if there are any numbers, then there must be an infinity of them. So, unlike the property of being a sibling, the property of being a number is not just anatomic but also holistic.

Part of coming to see why there must be lots of numbers if there are any is coming to see that being a number is really a relational property. (It’s evident on the face of it that being a sibling is a relational property; one speaks not only of *being* a sibling but also of *having* one.) Not all relational properties, however, are anatomic; a fortiori, not all relational properties are holistic. You can’t be a cat owner unless there is a cat that you own, so being a cat owner is a relational property. But it’s by no means obvious that you can’t be a cat owner unless there are other cat owners. Patently, the cat you own needn’t itself own a cat in order for *you* to own *it*. So, not very surprisingly, the relation between a cat owner and his cat is quite unlike the relation between a number and its successor; although being a cat owner and being a natural number are both relational properties, the latter is anatomic and holistic and the former is neither.

Or consider: you can’t earn the average income unless there are people whose incomes are related in a certain way to yours; so earning the average income is a relational property. In fact, it is the relation a wage and its earner bear to *n* wage earners and their wages if and only if (iff) his wage equals the sum of their wages divided by *n*. But there don’t have to be *other* people who

earn the average income in order for *you* to do so, so earning the average income is a punctate property. For that matter, there don't have to be *other* people who earn wages in order for you to earn the average income, since the relation that defines the average income is one that a wage earner can bear *to himself*. You can therefore earn the average wage in an economy in which you are the *only* wage earner. Indeed, in that sort of economy, you *can't but* earn the average wage.

Well, so be it. But why does any of this matter?

ANATOMISM AND THE THEORY OF LANGUAGE

Steamy philosophical issues can sometimes be rephrased as questions about whether a certain property is anatomic. This is a sort of ontological equivalent of the tactic of "semantic ascent"; and, like semantic ascent, it can have the salutary effect of lowering the temperature. Consider, for example, the steamy question of whether there could be a private language. To argue that there could be is at least to deny that the property of *having a language* is anatomic; correspondingly, it's at least to assert the conceptual possibility of a language with only one speaker. This doesn't get you *much* further, but it does help a little to distinguish the part of the private language problem that's about language from the part that's about privacy. (Compare Rhee's, "Can there be a private language?")

As previously remarked, this book is largely about whether semantic properties are holistic. We'll see that the standard argument for meaning holism requires the premise that semantic properties are typically anatomic. Discussing anatomism is thus a way into considering whether the connection between *being a symbol* and *belonging to a language* is internal; whether symbols can have their being only as parts of whole language systems. Since this will be our main expository tactic, we wish to alert the reader to a couple of caveats.

First caveat

Though questions about meaning holism can often be phrased as questions about whether some semantic property is anatomic, not just any semantic property will do for these purposes. This is just an uninteresting consequence of how "anatomic" was defined. So, for example, to claim that the property of *expressing the proposition that the cat is on the mat in L* is anatomic would be to claim that a language that has one expression that means that the cat is on the mat must also contain at least one other expression *that also means that the cat is on the mat*. This claim is most implausible; and, anyhow, no general issues about meaning holism would appear to turn on it.

The interesting and, *prima facie*, plausible claim is that *generic* semantic properties – loosely speaking, properties whose specification can be taken to involve variables ranging over propositions, contents, meanings, and the like – are anatomic. Examples of this claim are that *the property of expressing some proposition or other* or *the property of having some referent or other* or *the property of having some content or other* are anatomic. In particular, we'll see that it is for these sorts of properties that there is arguably an inference from semantic anatomism to semantic holism, so that if the first can be established, so too, perhaps, can the second. The reader is hereby advised that, barring specific notice to the contrary, when we talk about whether semantic properties are anatomic, it will almost always be generic semantic properties that we have in mind.

In particular, much of our discussion will be concerned with one of the following two, closely related doctrines. What we will call content holism is the claim that properties like *having content* are holistic in the sense that no expression in a language can have them unless many other (nonsynonymous) expressions in that language have them too. In effect, it's the doctrine that there can be no punctate languages. What we will call

translation holism is the claim that properties like *meaning the same as some formula or another of L* are holistic in the sense that nothing can translate a formula of L unless it belongs to a language containing many (nonsynonymous) formulas that translate formulas of L.³ It came as a surprise to us, and we hope it will interest the reader, to discover that almost all the arguments for meaning holism that actually get proposed in the literature are arguments for *content* holism. The argument for *translation* holism seems to be one that assumes that meanings supervene on intersentential relations – that they are something like inferential roles – and hence that translation preserves meaning only if the inferential relations among many of the sentences in the home language preserve the inferential relations among many of the sentences in the target language. We'll consider this sort of argument in detail in chapter 6.

Second caveat

The issues about anatomism aren't by any means the only ones that philosophers have had in mind when they raise "the" meaning holism question. For example, there's the thesis, famously explored by Wittgenstein, Austin, and their many followers, that there is an internal connection between *being a symbol* and playing a role in a system of *nonlinguistic* conventions, practices, rituals, and performances – an internal connection, as one says, between symbols and Forms of Life.⁴

We mention this, as it were, *anthropological* holism only to put it to one side. Our excuse for doing so is as follows. Anthropological holism is distinct from semantic holism *only* insofar as it concerns the relation between language and its intentional background – that is, the relation between language and the cultural background of beliefs, institutions, practices, conventions, and so forth upon which, according to anthropological holists, language is ontologically dependent. When applied to the background itself, however, anthropological holism just reduces to semantic holism. That is, it reduces to the doctrine that intentional states, institutions, practices, and the

like are ontologically dependent on one another; hence that they are anatomic. To put the point slightly differently, we have, at least for present purposes, no argument with the philosopher who holds that the linguistic is holistically dependent on an intentional background *but accepts atomism about the background*, thereby allowing, in effect, that there could be arbitrarily punctate Forms of Life. (We suppose, for example, that someone who is a "Gricean" about the relation between thought and language could coherently be an atomist about thought itself.) To put it yet another way, it might be that for anything *linguistic* to have content, there must be something *nonlinguistic* that has content. That's alright with us as long as the conditions for the nonlinguistic thing having content are atomistic.

Though anatomism isn't the only philosophical issue about semantic holism, it nevertheless suffices to distinguish two great traditions in the philosophy of language. The atomistic tradition proceeds from the likes of the British empiricists, via such of the pragmatists as Peirce and James. The *locus classicus* is the work of the Vienna Circle, but see also the Russell of *The Analysis of Mind*. The contemporary representatives of this tradition are mostly model theorists, behaviorists, and informational semanticists. Whereas people in this tradition think that the semantic properties of a symbol are determined solely by its relations to things in the nonlinguistic world, people in the second tradition think that the semantic properties of a symbol are determined, at least in part, by its role in a language. Languages are, inter alia, *collections* of symbols; so, if what a symbol means is determined by its role in a language, the property of *being a symbol* is anatomic. This second tradition proceeds from the likes of the structuralists in linguistics and the Fregeans in philosophy.⁵ Its contemporary representatives are legion. They include Quine, Davidson, Lewis, Dennett, Block, Devitt, Putnam, Rorty, and Sellars among philosophers; *almost* everybody in AI and cognitive psychology; and it may be that they include absolutely everybody who writes literary criticism in French.

It's pretty clear that whether semantic properties are anatomic is an interesting question if you happen to be interested in the philosophy of language. The point, to repeat, is that there is a widely (if often implicitly) endorsed argument which suggests that if a semantic property is anatomic, then it is also holistic. Suppose we grant, for the moment, that this inference from anatomism to holism goes through. Then anatomism about semantic properties has whatever consequences meaning holism itself has. And, arguably, the implications of meaning holism for the philosophy of language are formidable.

Dummett, for example, maintains that:

A thoroughgoing holism, while it may provide an abstractly intelligible model of language, fails to give a credible account either of how we use language as an instrument of communication, or of how we acquire a mastery of language. . . . The situation is essentially similar to that of a language all of whose sentences consist of single words, i.e. have no internal semantic structure; . . . it becomes unintelligible how the speakers of the language could ever have come to associate . . . senses with their unitary sentences, let alone to achieve the same association among different individual speakers; or how any one individual could discover the sense attached by another to a sentence, or decide whether it was or was not the same as that which he attached to it. In the same way, if a total theory is represented as indecomposable into significant parts, then we cannot derive its significance from its internal structure, since it has none; and we have nothing else from which we may derive it. (*Frege: Philosophy of Language*, pp. 599–600)

Dummett is, in effect, arguing from the following analogy: Sentences are interpersonally intelligible because their meanings are compositionally derived from those of their constituents and because speaker and hearer are privy to the meanings of the constituents and to the conventions that govern the derivation. This explanation presupposes that the constituents of sentences

are meaningful – indeed, that they mean the same in the speaker's language as they do in the hearer's. Similarly, Dummett claims, if I can understand your theory (by any incremental procedure), that must be because the content of the theory is determined by the contents of its constituent sentences. (Let's assume, for expository convenience, that theories are sets of sentences.) And if I can *learn* your theory (incrementally), that must be because I can learn part of your theory by learning some of its constituent sentences, more of your theory by learning more of its constituent sentences, and all of your theory by learning all of its constituent sentences. But, again, these possibilities presuppose that the sentential constituents of a theory *have* meanings – indeed, that they can have the same meanings in your whole theory and in the approximations to your whole theory that I learn along the way.

All of this would *seem* to be false if meaning holism is true, since, as the reader will recall, meaning holism would require that if any one sentence in your theory occurs in my theory, then practically all the sentences that occur in your theory must occur in my theory. And similarly, *mutatis mutandis*, if "theory" is replaced by "language." If holism is true, then I can't understand any of your language unless I can understand practically all of it. But then how, save in a single spasm of seamless cognition, could any language ever be learned?

We don't wish to take a stand on whether the considerations that Dummett advances constitute a refutation of semantic holism. For one thing, occasional digressions to the contrary notwithstanding, our business in this book is not to determine whether holism is true, but only to examine the arguments that have been offered in its favor. Second, suppose Dummett is right: suppose, that is, that the standard picture of how they are learned, communicated, and so forth presupposes that the semantic properties of theories and languages are determined by the semantics of their constituent sentences in something like the way that the semantics of a sentence is itself determined by the meaning of its constituent terms. Still, offering this

argument as an objection to meaning holism may underestimate the extent to which holists are likely also to be revisionists. A semantic holist might accept Dummett's analysis and reply, "So much the worse for our conventional understanding of how languages and theories are learned and communicated." Clearly Quine, Dennett, Stich, the Churchlands, and many other meaning holists are strongly tempted by this sort of revisionism.

Suffice it for present purposes that *if* you assume that properties like *having a meaning in L* and *having the same meaning as some expression in L* and the like are holistic, then a certain standard picture of how communication and language learning work would seem to be in jeopardy. The picture is that the linguistic and theoretical commitments of speaker and hearer can overlap *partially* to any degree you like: you can believe some of what I believe without believing all of it; you can understand part of my language without having learned the rest of it; and so forth. This would seem to be essential to reconciling the idea that languages have an interpersonal, social existence with the patent truth that no two speakers of the same language ever speak exactly the same dialect of that language. As Frege remarks in a related context:

Both the nominatum and the sense of a sign must be distinguished from the associated image . . . the image is subjective, the image of one person is not that of another. . . . [Hence] the image thereby differs essentially from the connotation [that is, sense] of a sign, which latter may well be a common property of many and is therefore not a part or mode of a single person's mind; for it cannot well be denied that mankind possesses a common treasure of thoughts which is transmitted from generation to generation. (Frege, "On sense and reference," pp. 159–60)

But, if we understand Dummett correctly, he is arguing that this picture of language as public property can make sense only to the extent that *partial* consensus in usage does *not* require *perfect* consensus of usage – that is, only to the extent that semantic holism is denied.⁶

So much for a first sketch of how issues about semantic anatomism may connect with some other questions proprietary to the philosophy of language.

BROADER IMPLICATIONS

Reference Holism and Scientific Realism

It has recently become increasingly clear that semantic holism also has repercussions further afield. Consider the property R that a linguistic expression has iff it refers to the same thing that some expression in English does. So, for example, R is a property that "la plume de ma tante" has (because it refers to the same thing that the expression "my aunt's pen" refers to), and so too do "la penna di mia zia" and, of course, "my aunt's pen." Question: Is the property R holistic? Could languages that overlap only slightly share any of their "ontological commitments?"

Here's one reason why this question matters. Suppose that ontological commitments are holistic, so that two languages can share any of their ontology only if they share quite a lot of it. It might then turn out, for example, that no language could have an expression that refers to what the English expression "the pen of my aunt" refers to unless it also has expressions that refer to, as it might be, Chicago, the cat's being on the mat, the last game of the 1927 World Series, the day after they built the Statue of Liberty, the last of the Mohicans, *The Last of the Mohicans*, and so forth.⁷ Such a result, though still of primary relevance to the philosophy of language, would nevertheless be interesting and rather strikingly counter-intuitive.

It raises the stakes, however, that the same considerations would apply if we asked about the semantic property R*. An expression has R* iff it refers to something or other that currently accepted astronomical theories refer to. Suppose that R* is anatomic, hence holistic on the assumption that anatomism implies holism. Then it might turn out that no

theory could refer to (for example) stars unless it could also refer to (as it might be) planets, nebulas, black holes, the center of the galaxy, the speed of propagation of light, and the location of the nearest quasar. It would follow that Greek astronomy (hence, Greek astronomers) couldn't ever have referred to stars. And it would follow from *that* that (what one had naively supposed to be) the Greek view that stars are very nearby and that they ride around the heavens on glass spheres is actually *not contested* by our view that the stars are very far away and don't ride around the heavens at all. In fact, strictly speaking, it would follow that the Greeks didn't *have* any views about stars; we can't, in the vocabulary of contemporary astronomy, say what, if anything, Greek astronomy was about. A fortiori, it makes no sense to speak of an empirically motivated choice between Greek astronomy and ours; whereof you cannot speak, thereof you must be silent.⁸

So if the property R* is holistic, then it may well turn out that scientific theories are empirically incommensurable unless their ontological commitments are more or less identical. But notice that *the* argument for Scientific Realism is that science is progressive; in the present case, the main argument for being Realistic about our astronomical theories is that, in virtue of having embraced them, we are in a position to make more and better predictions about stars than the Greeks did. If, as now threatens, it turns out that this is *trivially* true (because Greek astronomy made no predictions about stars at all or, indeed, about anything that *our* astronomy talks about), the standard argument for Scientific Realism goes down the drain.⁹

This understanding of the implications of R*'s being anatomic is widely shared. Inferring from holism about ontological commitment to anti-Realism (or relativism or Instrumentalism) about the theoretical constructs of science has been a main tactic of twentieth-century metaphysicians. Consider, among current practitioners, Quine, Goodman, Kuhn, Feyerabend, Putnam, and many others.¹⁰ Indeed, Kuhn's (putative) discovery of the incommensurability of scientific paradigms appears to be

the only result in recent philosophy that many nonphilosophers care about. And the argument that leads first to holism and then to incommensurability depends essentially on the claim that properties like R* are holistic (a fortiori, that they are anatomic).

Meaning Holism and Intentional Explanation

Now consider the property T. An expression has T iff it *translates* some or other expression of English. So, "the pen of my aunt" and "la plume de ma tante" have T, and so too do "The pen of my aunt is on the table" and "La plume de ma tante est sur la table," "La penna di mia zia é sul tavolo," and so forth. Question: Is the property T anatomic?

Here is why *this* question matters. Suppose, once again, that there is an argument from the anatomism of a semantic property to its holism. Then it might turn out that no language can have an expression that means what "The pen of my aunt is on the table" means unless it also has expressions that mean what, as it might be, "Two is a prime number," "London Bridge is actually in Arizona," "XYZ is not H₂O," "Snow is white," and "The snark is a boojum" mean. A consequence of this would be that Chaucer, Shakespeare, and Lincoln (and, for that matter, since XYZ is quite a recent invention, Ludwig Wittgenstein) did not speak a language in which one could say that the pen of one's aunt is on the table.

If there are arguments that show that neither Chaucer, Shakespeare, nor Wittgenstein could have said of his aunt's pen that it was on the table, then presumably much the same arguments would show that none of them could ever have *thought* of his aunt's pen that it was on the table. For, consider the property T* which a belief has iff it expresses a proposition that is the content of some belief of mine. According to the present assumptions, if T* is anatomic, then it is holistic. And if T* is holistic, then (assuming that thoughts are individuated by their propositional contents) it might turn out that nobody has thoughts that are tokens of the same type as my thought about

Auntie's pen unless he also has thoughts that are tokens of the same type as, as it might be, my thought that the cat is on the mat, my thought that black holes are odd kinds of objects, my thought that some presidents are wimps, or my thought that *Salome* will never sell in Omaha. This too might be considered an interesting, even counter-intuitive, result in the philosophy of mind. And once again there are implications further afield.

Lots of people, including most cognitive scientists and many riders on the Clapham omnibus, hold the following view of behavior: that higher animals act out of their beliefs and desires. According to this view, there are counterfactual-supporting generalizations that connect the mental states of higher animals with their behaviors (and with one another) *and which subsume mental states in virtue of their intentional contents*. Consider such shopworn examples as "If you see the moon *as being on the horizon*, then you will see it *as oversized*" or "If someone asks you what's the first thing salt makes you think of, you'll think of pepper" or "If someone asks you what's the first color you think of, you'll think of red."¹¹

And so forth.

We emphasize that it's in virtue of *what they are thoughts about* that thoughts fall under a generalization like "If you think of a color, the first color you think of is red" – that is, it's in virtue of their being thoughts about color and thoughts about red (reading "thoughts about" *de dicto*). A fortiori, the generalization subsumes you and me (as it might be) only if we both *have* thoughts about color and about red.

But now suppose that holism is true about thought content. Then, since you and I surely have widely different belief systems (think of all the things you know that I don't) and since, by definition, a property is holistic only if nothing has it unless many other things do, it may well turn out that none of your thoughts has the property of bearing T* to any of mine.¹² It would follow that not more than one of us ever has thoughts about color or thoughts about red. So, at most one of us is subsumed by the generalization that if you think of a color, then

the first color you think about is red. In fact, it might well turn out that, at most, one *time slice* of one of us is subsumed by this generalization since, after all, vastly many of one's beliefs change from moment to moment, and, on the present assumptions, belief individuation is holistic.

These sorts of considerations suggest that it might turn out that if T* is holistic, there are no *robust*, counterfactual-supporting intentional generalizations,¹³ none that is ever satisfied by more than an individual at an instant. Many philosophers have indeed drawn this sort of inference. Since, they argue, mental properties are holistic, there couldn't really be intentional laws; and since there can't really be intentional laws, intentional explanations can't be fully factual. (See, for example, Quine, Davidson,¹⁴ Stich, Dennett, both Churchlands, and others.) Presumably, if there aren't fully factual intentional generalizations, then there can't be an intentional science of human nature (or a scientific epistemology or a scientific moral psychology) in anything like the sense of "science" that the physical and biological sciences have in mind. "Behavioral science," "social science," "cognitive science," and the like are therefore, strictly speaking, oxymorons if semantic holism is true.

Above all, there can't be a scientific theory of *rationality*:

There are powerful universal laws obeyed by all instances of gold . . . but what are the chances that we can find powerful universal generalizations obeyed by all instances of rationally justified belief? The very same considerations that defeated the program of inductive logic, the need for a criterion of "projectibility" or a "prior probability metric" which is "reasonable" by a standard of reasonableness which seems both topic-dependent and interest-relative, suggests that . . . even in a restricted domain, for example physics, nothing like precise laws which will decide what is and is not a reasonable inference or a justified belief are to be hoped for . . . We should and must proceed in a way analogous to the way we proceed in science . . . ; but we cannot reasonably expect that *all* determined researchers are destined to

converge to one moral theory or one conception of reality. (Putnam, "Philosophers and human understanding," pp. 201–2)

Notice that this line of argument doesn't depend on parochial considerations about what you think intentional content *is*. All that's required is that, whatever it is, it's "topic-dependent and interest-relative" – namely, holistic.

The Autonomy of the Intentional

Our point up till now has been that the implications of meaning holism may reach far enough to jeopardize, on the one hand, a certain sort of Metaphysical Realism in the philosophy of science and, on the other hand, the likelihood that the intentional sciences might eventually produce theories whose objectivity and reliability parallel those of the physical and biological sciences. Prima facie, this makes meaning holism look like bad news from the point of view of linguists, psychologists, economists, cognitive scientists, and the like.

But, there is a more cheerful way of reading the moral; if the "constitutive principles" of intentional theories are ipso facto holistic (or normative or, maybe, holistic *because* normative; see chapter 5) in a way that those of the physical and biological sciences are not, then it may be that intentional explanations are ipso facto immune to a kind of reductive criticism with which the physical and biological sciences have sometimes seemed to threaten them. To put it the other way around, if you think of commonsense belief/desire psychology as "just another empirical theory," less articulate than, but not different in kind from, such philosophically unproblematic empirical theories as meteorology or geology, then it presumably follows that commonsense belief/desire psychology could turn out, on simply empirical grounds, to be largely or entirely *false* – just as it could turn out that much or all of our current meteorology or geology is simply empirically false. Commonsense belief/desire psychology *will* have turned out to be empirically false if, for

example, it proves not to be capable of integration with the rest of our developing scientific world view. The Churchlands (and maybe Quine) think that something like this is actually in the process of happening.

But that *couldn't* happen if, in virtue of their holistic character, interpretive and hermeneutic explanations are ipso facto not in competition with theories in the empirical sciences. One might then rationally take the view that the general structure of intentional explanation is, as one says in Britain, "not negotiable" *however* biology and physics turn out. There is in this line of argument more than a hint of the Kantian strategy of buying the autonomy of the foundations of ethics at the price of accepting a priori bounds on the scope of scientific understanding – except that it is now the conception of persons as intentional systems rather than the conception of persons as moral agents whose freedom from empirical critique is to be guaranteed by transcendental argument.¹⁵

Whichever way you look at it, if it's true that meaning holism is incompatible with a robust notion of content identity, and hence with a robust notion of intentional law, then the connections between the holism issues and some very deep questions about our understanding of ourselves are seen to be intimate and urgent.

CONTENT IDENTITY AND CONTENT SIMILARITY

Why, then, aren't many people outside philosophy (many cognitive/behavioral/social scientists, for example) worried about the holism issues? One reason is that they may not have noticed the undesirable consequences of holism, or they may doubt that these consequences actually follow. Another reason is that it's widely supposed that even if holism precludes a robust notion of content *identity*, still it permits a robust notion of content similarity. (There's a third reason too, as we'll see in the next section.) Taking this for granted seems like just

common sense. After all, there does *seem* to be a colloquial notion of belief similarity. We do say things like “What I believe is a lot like what the President believes” or “Her world view is sort of similar to Dracula’s” or “His understanding of definite descriptions is less like Russell’s than like Strawson’s” and so on. So maybe this colloquial sense of “similar belief” can be co-opted to provide for a robust formulation of intentional generalizations. Maybe the right generalization is: If somebody asks you something *sort of like* what is the first color you think of, then you will think of something *sort of like* red.¹⁶

The trouble is that we really have no idea what it would be like for this new generalization to be true (or false) and, barring some illumination in this quarter, the suggestion that appealing to content similarity may mitigate the severer consequences of semantic holism is simply empty. This point is so important, and so widely goes unrecognized, that we propose to spend a little time rubbing it in.

No doubt, one does know (sort of) what it is like to more or less believe the same things as the President does; it’s to share *many of the President’s beliefs*. For example, the President believes P, Q, R, and S, and I believe P, Q, and R; so my beliefs are similar to his. An alternative, compatible reading is: the President believes P and Q *very strongly* and I believe them equally strongly or almost as strongly, so again my beliefs are similar to his. But neither of these ways of construing belief similarity helps with the present problem. The present problem is not to make sense of believing-most-of-P, -Q, -R, -and-S or of more-or-less-strongly-believing-P; it’s to make sense of believing *something-similar-to-P* – that is, *more-or-less-P*.

The colloquial senses of “similar belief” *presuppose* some way of *counting* beliefs, so they presuppose some notion of belief *identity*. If you have most of the beliefs that I have, then, a fortiori, there are (one or more) beliefs that we both have. And if there is a proposition that you sort of believe and that I

believe strongly, then, a fortiori, there is a proposition that is the object of both of our beliefs. But precisely because these colloquial senses of belief similarity *presuppose* a notion of belief *identity*, they don’t allow us to *dispense with* a notion of belief identity *in favor of* a notion of belief similarity. In consequence, if you’re a holist and your notion of belief identity is very unrobust, so that, de facto, people can hardly ever have the same belief, then it will also turn out that, in either of the colloquial senses just discussed, people can hardly ever have *similar* beliefs. If it’s never true that I believe *any* of what the President believes, then, of course, it can’t be true that I ever believe *most* of what he believes. If the President and I never believe the same thing, then there is nothing that he believes as strongly as I do.

It’s not, of course, incoherent to imagine a notion of “similar belief” which, unlike these colloquial ones, is compatible both with meaning holism and with there being robust intentional generalizations. The trouble is, as we remarked above, that nobody seems to have any idea what this useful new sense of “similar belief” might be. On the contrary, it seems sort of plausible that you can’t have a robust notion of *similar* such and suches unless you have a correspondingly robust notion of *identical* such and suches. The problem isn’t, notice, that if holism is true, then the conditions for belief identity are hard to meet; it’s that, if holism is true, then the notion of “tokens of the same belief type” is defined *only* for the case in which *every* belief is shared. Holism provides no notion of belief-type identity that is defined for any other case and no hint of how to construct one. But if there is no construal of the claim that two beliefs are tokens of the *same* type in cases where belief systems fail to overlap completely, how, in such cases, are we to construe the notion of two beliefs being tokens of *almost* the same type? (One recent proposal for construing the notion of similarity of meaning will be discussed in chapter 7, q.v.)

We really do think it’s hard to get out of this; the sort of unconsidered talk about similarity of intentional content that is

currently so prevalent in cognitive science serves only to obscure the magnitude of the problem. For example, it might be suggested that a content holist could endorse a *physicalistic* account of belief similarity; after all, your beliefs and mine are presumably *identical* if you and I are *identical* molecule for molecule.¹⁷ Doesn't it follow that our beliefs are *similar* if we are *similar* molecule for molecule? This notion of belief similarity would be robust because, even if no two time slices of organism are ever physically identical, there are plenty of ways, surely, that two time slices of organism can perfectly well be physically *alike*.

But, on second thought, this doesn't help at all. Even if it's granted that identity of belief systems supervenes on physical identity, it doesn't begin to follow that similarity of belief systems supervenes on physical similarity. It is, perhaps, reasonable to assume that if you are my molecular twin, then you share all my beliefs. But it is entirely gratuitous to assume that if you are my molecular cousin, some of your beliefs are ipso facto similar to some of mine. (Which ones, by the way?) No doubt there *are* indefinitely many ways in which the brains of molecular cousins are similar; but there are also *indefinitely many ways in which they aren't*, and we have no idea how to decide which similarities and differences are the ones that determine whether their beliefs are similar. Which is just to say that nobody has a better idea of how to explicate a notion of *physical* similarity that is relevant to psychological taxonomy than of how to explicate a notion of *content* similarity that is relevant to psychological taxonomy.

"Well, maybe two beliefs are similar if they participate in mostly the same inferences." There are two reasons why this too doesn't help. One is the same sort of point we've just been noticing: that if this proposal is to provide a robust notion of similar belief, it will have to presuppose a correspondingly robust notion of identity of inference; and that is one of the things that meaning holism appears likely to deny us. If it turned out to be a consequence of meaning holism that no two

people ever have exactly the same belief, it would surely also turn out to be a consequence of meaning holism that no two people ever accept exactly the same inference. After all, *identical* inferences must have *identical* premises and *identical* conclusions. And if it is replied that, well, holism still allows that different people could accept *similar* inferences, we're back where we started – except that it's the notion of *similar inference* rather than the notion of *similar belief* that now cries out for explication.

The second problem with reconstructing similarity-of-beliefs-entertained by reference to similarity-of-inferences-endorsed is that some inferences have to count for more than others, surely. Consider the man who may be thinking about red. When I think about red, I am in a state from which I am prepared to make certain inferences about tomatoes. So, for example, if I think this book is red, then I'm prepared to believe that this book is the same color, more or less, as ripe tomatoes are. But my willingness to make *this* inference (and thousands like it) surely can't be *constitutive* of my having thoughts about red. If it were, Shakespeare would be out of luck; he didn't know about tomatoes.

In fact, however you individuate beliefs, it's sure to turn out that there are *vast numbers* of red things – hence vast numbers of things about red – that I know about that Shakespeare didn't; and, of course, vice versa. So now we need to know *how much* the differences between the red-inferences I endorse and the ones that Shakespeare did count as differences in our concept of red. The extent to which this sort of question lacks a principled answer is the extent to which we have no notion of similarity of content that is compatible with a holistic account of belief attribution. And it lacks a principled answer entirely; does believing that Mars is red count more or less for having the concept *red* than believing that tomatoes are?¹⁸

The long and the short of it seems to be that intentional explanation needs a robust notion of belief identity, and meaning holism appears to prejudice the possibility of such a

notion. You can't get out of this just by appealing to a notion of *similarity of content*, because all the robust notions of content similarity – or, at a minimum, all the ones that spring to mind – *presuppose* a robust notion of belief identity and hence are themselves incompatible with holism if robust belief identity is.

MEANING HOLISM AND THE ANALYTIC/ SYNTHETIC DISTINCTION

There is an alternative move that it's traditional for philosophers to make at this point – namely, to opt for a notion of belief identity after all, one that's grounded in an analytic/synthetic distinction. Beliefs are identical iff they participate in the same *analytic* inferences. (Presumably a corresponding notion of belief similarity can be introduced if it's required, some variant on "Beliefs are similar insofar as they participate in *many of the same* analytic inferences." See chapter 2 for further discussion of analyticity and chapter 6 for its relation to belief identity.) *Strictly speaking*, this way of squaring content holism with a robust notion of belief similarity might surely be accused of begging the question, since, once again, it appears that a robust notion of *accepting the same inference* (hence a robust notion of *same inference*) is being taken for granted. But we propose not to harp on this any longer. The a/s distinction has been lurking in the closet through this whole discussion, and it is now time to let it out.

Up till now, we've been considering some consequences of assuming both that semantic properties are typically anatomic and that if a semantic property is anatomic, then it is holistic. Notice that the first assumption is relatively innocuous unless the second one is also in place. It would, no doubt, be interesting and curious to show that, for example, you can't share any of my beliefs unless you share at least two of them (*mutatis mutandis*, that a language can't express any propositions unless it can express several, and so forth). But it's not at

all obvious that drastic implications for theory commensurability, Scientific Realism, translation, intentional explanation, and the like would follow from this sort of "molecularist" semantics.¹⁹ These seem to depend on the *holistic* claim that the conditions for content relativize to entire languages or belief systems; for example, that you can't share any of my beliefs unless you share practically all of them.

What we now want to emphasize is that the argument from anatomism to holism itself depends on the premise that no principled a/s distinction can be drawn. If this is so, then the only context in which a discussion of semantic holism is worth having is one in which the failure of the a/s distinction is taken as common ground. We remarked, in the preceding section, that if not many cognitive scientists are worried about the threat that holism poses to the concept of belief identity, that's often because they suppose that some notion of belief similarity will serve to take up the slack. In like spirit, if not many "functional role" semanticists or verificationists are disturbed by the spectre of holism, that's often because they are prepared to buy into some kind of a/s distinction.²⁰

We now propose to consider how the argument from the anatomism of semantic properties to the holism of semantic properties might be supposed to run and what role in the argument the denial of the a/s distinction plays. Here's a candidate formulation.

Argument A

Premise 1: Generic semantic properties like T, T*, R, R*, being-some-or-other-belief-of-Smith's, being a formula of language L, etc. are anatomic.

Comments:

1. We want to be noncommittal about *how many* generic semantic properties are anatomic. The argument under analysis

Centers
of
chapter

requires only that the property of being-some-or-other-belief-of-Smith's is.

2. Premise 1 might itself be derived from, for example, some version of "inferential role" semantics – for example, from the assumption that the identity of a concept (*mutatis mutandis*, the meaning of a word) is at least partially determined by its role in a belief system (or language or theory). Our impression is that most contemporary philosophers who accept premises like 1 do so for this sort of reason. For example, we'll see in chapter 2 that Quine is widely read as endorsing a form of argument A in which the first premise derives from verificationist assumptions about semantics: roughly, the assumption that the content of a belief is the means of determining that the belief is true (/false), including, in particular, the inferences involved in such determinations.

Lemma: If Smith has the belief that P, he must have other beliefs not identical to P.

Comment: Instantiation.

Premise 2: There is no *principled* distinction between the propositions that Smith has to believe to believe that P and the propositions that Smith doesn't have to believe to believe that P.

Comments:

1. The standard reason for holding premise 2 is that, on the one hand, the only principled distinction anyone can think of depends on the idea that if you can't believe P unless you believe Q, then "if P, then Q" must be analytic (or, perhaps, analytic *for you*), and, on the other hand, there is no principled *a/s* distinction. This, then, is the precise point at which the argument from the anatomism of semantic properties to the holism of semantic properties turns on the rejection of *a/s*.

2. Notice, however, that what is meant by "the rejection of *a/s*" in this context is quite different from the rejection of the *a/s* distinction that Quine almost certainly has in mind in "Two dogmas of empiricism" and other of his papers in which the distinction is impugned. When Quine says "No *a/s*," he presumably means "No analytic sentences." On that reading of premise 2, however, argument A would appear to be inconsistent. For premise 1 requires that there be sentences other than P that must be believed if P is believed; and it looks as though the hypothetical formed by writing one of these sentences after "if P, then" must be analytic, as we've just observed.

So the reading of premise 2 that argument A really requires is "The *a/s* distinction isn't principled," rather than "No sentences are analytic." This reading is, of course, acceptable to conceptual role semanticists and their ilk; but it's pretty clear that Quine couldn't put up with it.²¹ There is a fair amount of irony in this. Accepting semantic holism is often seen as a *consequence* of agreeing with Quine about the *a/s* distinction. But what Quine said about the *a/s* distinction is that there are no analytic sentences, and it doesn't look as if a semantic holist who endorses argument A *can* agree with *that*.

Conclusion: The property of being-some-or-other-belief-of-Smith's is holistic.

Comments:

1. The reference to Smith is inessential. If the argument is right, it shows that there couldn't be a punctate mind (a mind which can entertain only one proposition) or, *mutatis mutandis*, a punctate language (a language which can express only one proposition).

2. The form of argument A is: "If some *a*'s are F, and there is no principled difference between the *a*'s that are F's and the ones that aren't, then all *a*'s are F." So argument A has the form of a "sorites" or "slippery slope."

See
p274

3. As often happens when a form of philosophical argument is in the air, it's frustratingly difficult to find fully explicit instances in print. (Devitt, "Meaning holism," registers the same complaint.) Stich, however, comes pretty close:

I want to demonstrate that . . . intuitive judgments about whether a subject's belief can be characterized in a given way . . . are often very sensitive . . . to other beliefs that the subject(s) are assumed to have. The content we ascribe to a belief depends, more or less holistically, on the subject's entire network of related beliefs.

.
Consider the fact that . . . intuitions [of conceptual identity and difference] . . . seem . . . to lie along a continuum. Recall, for example, the case of Mrs. T, the woman suffering from gradual, progressive loss of memory. Before the onset of her illness Mrs. T clearly believed that McKinley was assassinated. By the time of the dialogue reported in Chapter 4 she clearly did not believe it. But at what point in the course of her illness did her belief stop being content-identical with mine? The question is a puzzling one and admits of no comfortable reply. What we are inclined to say is that her belief gradually becomes less and less content-identical with mine [as the inferences we share come to overlap less and less] . . . How much physics must my son know before it is appropriate to say that he believes that $E = MC^2$? The more the better, of course, but there are no natural lines to draw. (Stich, *From Folk Psychology to Cognitive Science: The Case Against Belief*, pp. 54, 85-6)

More or less explicit versions of argument A are also to be found in Dennett ("Intentional systems"), Churchland ("Perceptual plasticity and theoretical neutrality: a reply to Jerry Fodor"), Gibson (*The Philosophy of W. V. Quine: An Expository Essay*), and maybe in Quine ("Two dogmas"). Indeed, the practically universal tendency to invoke "No *a/s*" as a premise when making a case for semantic holism would seem senseless except in the context of some such argument as A.

The status of argument A

There are plenty of reasons for doubting that argument A is sound. In the first place, one might doubt that semantic properties actually are anatomic. For example, the usual justification for supposing that they are is that one assumes some functional role (or verificationistic) account of meaning. To deny this assumption would thus be to undermine the standard argument for premise 1. Second, even if the *a/s* distinction is untenable, there might be some other principled way of distinguishing the propositions that you have to believe to believe P from the ones you don't. Third, as we remarked above, A is a species of sorites argument, and these are notorious for leading from true premises to false conclusions. Consider the slippery slope that runs from there being no principled difference between baldness and hairiness to the conclusion that everyone is bald (or that nobody is). Fourth, A's validity might be challenged on grounds independent of the status of slippery slope arguments. At least one version of this fourth objection merits discussion.

*Weak anatomism, strong anatomism, and the *a/s* distinction*

A number of people (Boghossian, Loewer, Maudlin; see n. 9) have suggested the following as a situation in which A's premises might be true *but its conclusion false*. Imagine that there are disjoint sets of propositions such that (1) believing any one of these sets is sufficient for being able to believe P; (2) you must believe *at least* one of these sets in order to believe P; (3) none of these sets is such that you must believe *it* in order to believe P. So someone can believe P if he believes A *or* if he believes B, and so on. It might be further assumed that there are indefinitely many such sets of sufficient-but-not-necessary conditions for believing that P, and that nobody is able to believe the proposition formed by disjoining these indefinitely many sets of propositions (perhaps because the resulting disjunctive proposition is so complicated that no mind is able to

entertain it). Then, on the one hand, premise 1 would clearly be true. And premise 2 would be true in at least the sense that there are no analytic beliefs.²² Yet neither content holism nor translation holism would follow. Content holism wouldn't follow because it requires that there must be *many* other propositions that I believe if I believe that P (that is, it requires a *lot* of anatomism), and the current assumptions allow that some or all of the disjoint sets each of which is sufficient to be able to believe P might be quite small. So, compatible with the present account, content might be molecular rather than holistic. Translation holism wouldn't follow either, because it requires that for two people to share any belief, they must share at least one other belief, and the present model allows that what you believe is P and A, whereas what I believe is P and B. What everybody *really* wants is that meaning should be anatomic and that translation holism should nevertheless be false. This suggestion seems to do the trick.

There is, to put the point slightly differently, a quantifier-scope ambiguity lurking in the definition of "anatomic," and hence in premise 1 of argument A. What might be meant by claiming that properties like *having the belief that such and such* are anatomic – and what we have thus far been meaning in discussing the claim – is:

There are other propositions such that you can't believe P unless you believe them. Call this the "long scope" or "strong" reading.

Or what might be meant is:

You can't believe P unless there are other propositions that you believe. Call this the "short scope" or "weak" reading.

The proposed criticism of argument A is that, on the short scope reading of "anatomic," the premises can be true and the conclusion false.

The line of thought we are considering is framed as an objection to argument A. But it might equally be thought to show that holding anatomism while rejecting the a/s distinction need entail no pernicious consequences. As we've been seeing, on the short scope reading of "anatomic," conjoining premises 1 and 2 would *not* entail that you can't share any of my beliefs without sharing all of them, or even that your having any one belief requires your having lots of others. In short, it looked at first sight as though argument A might make semantic holism the only coherent alternative to semantic atomism. That is, it looked at first sight as though the only way to avoid argument A might be to take premise 1 to be *false*. But if, instead, you take premise 1 to be *true on the short scope reading*, then atomism is blocked, yet holism doesn't follow. Atomism is not conceded, but argument A is nevertheless defanged.

The trouble with this line of thought is that the kind of anatomism you get if you take premise 1 on the short scope reading is too weak to be worth the effort of defending. The way to see this is to ask yourself why it ever seemed important to argue that semantic properties are anatomic. We think that the answer is pretty clear: There is undeniably a pre-theoretic intuition that two people couldn't agree about *only one thing*. The intuition is that, if you and I agree that protons are very small, then there must be lots of other propositions we agree about too – for example, that protons aren't tangerines or prime numbers or mammals; that, *ceteris paribus*, very small things are smaller than very big ones, that there are sub-atomic particles, that positive charges are different from negative charges, and so forth. In effect, semantic holism proposes to hold onto this intuition even if the price is claiming that we can't agree that protons are very small unless we agree about *everything* else.

We're not prepared to endorse this intuition straightaway; to do so would just close the book against the possibility of semantic atomism.²³ But we don't deny its first blush force. One might even think that the very point of content attribution turns

upon the intuition being true; that it's only because we're guaranteed that people who share *any* beliefs must share *lots* of them that content attributions can warrant predictions "from the intentional stance." (In chapter 5 we'll examine an argument of Dennett's that's much in this spirit.) Our present point, however, is that if honoring this intuition is the motive for anatomism, then weak anatomism isn't any better off than atomism is.

The holist wants to capture the intuition that you and I can't both believe the proposition that protons are very small unless we also both believe some other propositions. But beware of the quantifier ambiguity here too. This might mean "Unless each of us believes at least one proposition other than '*Protons are very small,*'" or it might mean "Unless there is at least one proposition other than '*Protons are very small,*' that we both believe." It's clearly the second reading that is demanded by the idea that you and I couldn't agree on just one thing. (The first reading says only that neither of our beliefs that protons are very small could be punctate.) But the second reading is just *strong* anatomism; that is, it's not one to which a *weak* anatomist is entitled.

The sum and substance of this is that strong anatomism is the only kind worth having. So, from now on, we'll be understanding premise 1 according to the long scope interpretation.

The status of argument A (continued)

It's still on the cards, of course, that there may be something wrong with arguments that seek to infer semantic holism from anatomism together with the rejection of a principled a/s distinction. But though it's not clear what one should say about such arguments, the following *is* clear: If there is a principled a/s distinction, then the *inference* from anatomism to holism is blocked. A principled account of the a/s distinction would distinguish the propositions that you do have to believe to be able to believe that P from the ones that you don't (and the propositions that a language has to be able to express if it's able

to express P from the propositions that it doesn't). So far as we can see, this point is perfectly general. *Whatever* your argument for semantic holism might be – whether or not it's some version of argument A – it's going to fail if the a/s distinction can be sustained. It's only if you contemplate giving up an a/s distinction that you have to contemplate taking semantic holism seriously.

So, then, if we're going to discuss semantic holism at all, it had better be common ground that premise 2 of argument A is OK; and, specifically, that it is OK because there is no principled a/s distinction. Notice that, for our purposes, this rules out any possibility of a "molecularist" compromise between atomism and holism. A molecularist says that there are other beliefs that we must also share if we are to share the belief that P, but he denies that *all* our other beliefs have to be shared. But distinguishing between those that do and those that don't depends on invoking the a/s distinction, for believing P requires accepting the *analytic* inferences in which P figures. Molecularism is thus a closed option on the only assumption on which holism is sufficiently plausible to be worth discussing: namely, that the a/s distinction can't be sustained.

So be it. In what follows we will be seeking to undermine A and arguments like it; but (unlike Devitt, "Meaning holism," for example) we won't claim that what's wrong is that premise 2 is false. We also won't claim that what's wrong with arguments like A is that they are slippery slopes – though it may well be that all arguments from anatomism to holism *are* slippery slopes, which may well be one of the things that are wrong with them. What we'll do instead is attack the grounds that have been alleged in support of the *first* premise; we'll try to show that no good reason has yet been given as to why (generic) semantic properties are (strongly; see above) anatomic. If we are right about this, then, a fortiori, there are no good arguments for semantic holism, it being the *stronger* thesis.

One last preliminary remark about argument A: though it isn't, perhaps, really very convincing, it may nevertheless be a straw in the wind. At a minimum, if you are *independently*

persuaded that the atomistic alternatives to holism have been explored *and have been shown not to work*, then arguments in the spirit of A may well suffice to produce a rational conviction that holism is true. That brings us to the next part of our story.

SEMANTIC ATOMISM

Why is almost everyone a meaning holist? There are, we think, two kinds of considerations conducive to the doctrine. The first consists of positive arguments (in the spirit of argument A, for example) that meaning holism is true. The second is a sort of intuition about the historical situation in semantics. It's the intuition that holism is the last log afloat, that the history of philosophical discussions of meaning shows that either semantic properties are holistic or there are no such properties.

Suppose you think that there is no a/s distinction and that there is a convincing argument from atomism to holism. In consequence, you think that semantic properties must be either holistic or punctate. What is the likelihood that they are punctate? Well, if they are, then, by definition, the meaning of an expression can *not* depend on its role in a language. What else might it depend upon? The traditional nonholist answer is: some symbol/world relation – specifically, some punctate symbol/world relation, some relation that one thing could bear to the world even if nothing else did. This is the doctrine we've been calling “semantic atomism.”

It's a widely held view that much of the history of the philosophy of language consists of a failed attempt to make semantic atomism work.²⁴ Given this view, there is an inductive argument that the only story about language that is compatible with taking semantic properties seriously is holism.²⁵ For example, the tradition that runs from the mentalistic empiricism of Hobbes, Locke, Berkeley, and Hume to the behavioristic empiricism of Watson, Mead, Skinner, Dewey, and Ogden and

Richards offers two different reconstructions of the mind/world relation on which content is supposed to depend. Both of these reconstructions are atomistic, and both of them fail.

The mentalistic version of this tradition holds that semantic properties inhere, in the first instance, in a certain class of mental particulars, in “Ideas” according to one use of that term. (The semantic properties of, for example, English words are derivative; to have a word that means *dog* is to have a word that is associated, in the right way, with the *dog* Idea.) These mental particulars are species of images, and what they mean depends on what they resemble.²⁶ To have the idea of a dog is thus (approximately) to have an Idea that looks like a dog; to have the idea of a triangle is (approximately) to have an Idea that looks like a triangle. And so forth. Since what one of one's Ideas looks like is presumably independent of what other Ideas one has, the requirements for meaning are atomistic according to this account.

According to the behavioristic version of the tradition, meaning inheres in the first instance in certain (paradigmatically verbal) behavioral gestures. To have in one's behavioral repertoire a sound that means *dog* is (approximately) to be so conditioned that dogs reliably cause one to utter that sound; to have in one's behavioral repertoire a sound that means *triangle* is (approximately) to be so conditioned that triangles reliably cause one to utter *that* sound. And so forth. Since whether one's behavioral repertoire includes a sound the utterance of which is reliably conditioned to dogs is, presumably, independent of what, if anything, the other sounds in your repertoire are reliably conditioned to, the requirements for meaning are atomistic according to this account too.

We propose to spare the reader a rehearsal of the arguments which show that meaning can't be reduced either to resemblance or to behavioral conditioning.²⁷ We remark only that to admit that these versions of meaning atomism are hopeless is not the same as admitting that meaning atomism is false; a fortiori, it's not the same as admitting that meaning holism is true. In fact –

sell

or so it seems to us – the present situation in the philosophy of language includes the following open options:

✓ 1. It might turn out that semantic properties are anatomic (so that semantic atomism is false) but that holism doesn't follow because the a/s distinction proves to be tenable. What would be left is a sort of semantic molecularism (as has been suggested by Dummett, among many others). Roughly, the smallest language that could express the proposition that P would be one that can express the propositions to which P is analytically connected.

✗ 2. It might turn out that semantic properties are anatomic (so that semantic atomism is false) but that holism doesn't follow because, although the a/s distinction isn't tenable, there is some other principled way of grounding the distinction between the inferential relations that are constitutive of content and the ones that aren't. Once again, the upshot would probably be some sort of semantic molecularism.

✗ 3. It might turn out that holism follows from the assumption that semantic properties are anatomic, but that semantic properties aren't, in fact, anatomic. That is, it might turn out that meaning atomism is true.

✗ What the familiar arguments show, it seems to us, is that, if option 3 is the way it does turn out, then somebody will have to cook up a story about how symbol/world relations are constitutive of content that does *not* appeal to resemblance or behavioral conditioning. What we doubt is that the reasons that have thus far been invoked against meaning atomism show that this *could not* happen.²⁸

Modesty, however, is our middle name; nothing so ambitious as a defense of meaning atomism is contemplated in the text that follows. Here is what we propose to do instead. We want to look, as carefully and exhaustively as we can, at arguments for meaning holism that reject *an als distinction* but that *do not assume that meaning atomism has*

been shown to be false. We're going to try to show that none of these residual arguments is convincing. The bottom line might then be that there aren't any semantic properties; or it might be that some kind of meaning atomism is true but nobody knows *which* kind; or it might be that there really are good arguments for meaning holism, but nobody has been able to find one yet. We're noncommittal; you choose.

W. V. O. QUINE:
MEANING HOLISM AND
CONFIRMATION HOLISM

Quine's "Two dogmas of empiricism" is perhaps the most analyzed short philosophical paper written in the last fifty years and is the point of departure for many discussions of holism. Our treatment will therefore be relatively brief. We propose to swallow Quine's epistemology whole, confining ourselves entirely to those aspects of "Two dogmas" that are widely interpreted as bearing on the status of meaning holism. We think that the implications of Quine's epistemological doctrines for strictly semantic theses have been in some respects exaggerated; in particular, that meaning holism can be resisted even assuming – as we are inclined to do – that most of what "Two dogmas" says about the holism of confirmation (the so-called Quine/Duhem (Q/D) thesis) is right-headed.

We commence by reminding the reader of the architecture of "Two dogmas." In effect, the paper divides into two major parts. The first part (sections 1–5) is a long argument to the conclusion that the analytic/synthetic distinction cannot be enforced, together with a rejection of the reductionist claim that there is, for each "statement" in an empirical theory, a proprietary range of (dis)confirming conditions (for example, a proprietary range of sense data statements). By definition, a reductionist holds that the confirmation conditions for a statement are knowable a priori, because they are among the statement's *analytic* implications. Quine thus takes anti-

reductionism to be an immediate consequence of abandoning the *a/s* distinction. To claim that the *a/s* distinction fails is to claim, in particular, that our knowledge of confirmation relations is a posteriori. This means, for Quine, that it is *contingent* knowledge, since, like the positivists before him, Quine assumes that "a posteriori" and "contingent" are co-extensive (compare Kripke, *Naming and Necessity*).

Except for the very last paragraph of section 5 (which will presently be the object of much of our discussion), the moral of this first part of "Two dogmas" is clear; the consequence of assuming that there is no *a/s* distinction is confirmation holism: "Our statements about the external world face the tribunal of sense experience not individually but only as a corporate body" ("Two dogmas," p. 41). This, of course, is the part of "Two dogmas" that has received most of the critical attention. But we don't propose to get involved either in the squabble about whether Quine really did *show* that the *a/s* distinction is untenable or in the squabble about whether the failure of the *a/s* distinction really implies the Q/D thesis that confirmation is holistic. Our strategy is rather to grant all this and to ask what the implications for semantics are.

The second (and much the shorter) part of "Two dogmas" (section 6) spells out what Quine takes to be the pragmatist implications of confirmation holism for the philosophy of science: "As an empiricist I continue to think of the conceptual scheme of science as a tool, ultimately, for predicting future experience in the light of past experience . . . in point of epistemological footing the physical objects and the gods differ only in degree and not in kind. Both sorts of entities enter our conception only as cultural posits" ("Two dogmas", p. 44). Whereas our official position is that we concede both the failure of the *a/s* distinction and the truth of confirmation holism, we are officially neutral about pragmatism; it turns out to be largely independent of the semantic issues that we want to raise.

* (It is worth a brief digression to emphasize that this is so; you can coherently buy Quine's confirmation holism *without*

buying his pragmatism; prima facie, they are independent in both directions.)

That a pragmatist might accept an *a/s* distinction (and hence be a localist about confirmation) is just a way of saying that a pragmatist can be what Quine calls a "reductionist" (that is, what is usually called a "verificationist"⁽¹⁾), which, of course, pragmatists typically *have* been (see Peirce, for example). Indeed, one of the things that's so *surprising* about "Two dogmas" is precisely the idea that you can have a *non-reductionist* pragmatism (that is, a pragmatism without an *a/s* distinction, hence a pragmatism that is holist about confirmation). Given the history of pragmatism, you might have supposed that its connection to reductionism is internal.

That a confirmation holist doesn't have to be a pragmatist – indeed, that he can be an ontological Realist⁽²⁾ – also seems clear. In fact, if the Q/D thesis is the claim that it is a posteriori (rather than a priori) what confirms what, then the Q/D thesis is a doctrine that a Realist clearly ought to embrace. The reason is this: from the Realist perspective, *what confirms what* is a matter, not of linguistic convention, but of what is actually (for example, causally) connected to what *in the world*. If the pinkness of the litmus confirms the acidity of the fluid, that is not because "is acid" *means* (as it might be) "turns litmus pink," but rather because *being acid* and *being a cause of pinkness in litmus paper* are lawfully connected properties of acids. But if confirmation relations depend on how things are in the world, then presumably our knowledge of confirmation relations must be a posteriori – hence the Q/D thesis.

Or, rather, hence the Q/D thesis according to one of its formulations.³ As will become increasingly clear, we think that "the" Q/D thesis is really a galaxy of nonequivalent (but closely interrelated) doctrines and that "Two dogmas" uses various versions of the thesis as its polemical purposes require. For example, we'll see that Quine has an important use for a formulation of the Q/D thesis according to which you can hold onto any statement, if confronted by recalcitrant data, by

making compensatory adjustments elsewhere in your theory. This isn't the same claim as that confirmation relations are ipso facto a posteriori, but it's hard to see how the first could be true unless the second were too. Since Quine assumes the equivalence of the a priori with the semantic, revising a *nonempirical* relation would presumably involve equivocation.

Since we think that most of the epistemological doctrines that cluster together in the Q/D galaxy are probably right, we will not insist on the differences between them. Sooner or later, somebody will doubtless write a doctoral dissertation that sorts them all out.

Squeezed in between the long first part and the short second part of "Two dogmas" is the following paragraph:

The idea of defining a symbol in use was, as remarked, an advance over the impossible term-by-term empiricism of Locke and Hume. The statement, rather than the term, came with Frege to be recognized as the unit accountable to an empiricist critique. But what I am now urging is that even in taking the statement as unit we have drawn our grid too finely. *The unit of empirical significance is the whole of science.* ("Two dogmas," p. 42; our emphasis)

If there is a holist semantic thesis in "Two dogmas," this would appear to be it.

We say "If there is" because it would be natural, given both its position in the text and the occurrence of phrases like "unit accountable to an *empiricist critique*" and "unit of *empirical significance*" (our emphasis), to wonder whether this passage isn't itself just a reiteration of the Q/D thesis – that is, of holism about *confirmation*. Certainly it's *at least* that. However, three considerations at least *suggest* that it's worth taking seriously the view that a semantic thesis is at issue: first, the reference to Frege⁴ (who, presumably, really was talking about the units of meaning rather than the units of confirmation); second, the critical tradition according to which "Two dogmas" is a *locus*

classicus for semantic holism,⁵ and third, since reductionism is explicitly viewed as both a semantic and an epistemological doctrine in "Two dogmas," it's natural to construe its denial there as both an epistemological and a semantic doctrine too. It is quite plausible (though, as we will see, by no means inevitable) to suggest that, just as "Two dogmas" offers the Q/D thesis to oppose (say) Carnap's localism about confirmation, so too it offers semantic holism to oppose Carnap's localism about meaning.

We propose, then, for purposes of argument, to just drop the "empirical" and read Quine as claiming that "the unit of *significance* is the whole of science" (*mutatis mutandis*, the whole of a belief system). The question then arises as to how this claim is related to the text it immediately follows: that is, "What does confirmation holism have to do with semantic holism?" to put this question in a nutshell.

There is an obvious suggestion – one which we take to be congenial to the critical tradition in Quine scholarship. (See, for example, Dummett, *Frege: Philosophy of Language*; Gibson, *The Philosophy of W. V. Quine*.) Quine is a verificationist,⁶ that is, he accepts the identification of the meaning of a statement with its means of confirmation; as Quine puts it, paraphrasing Peirce, "The meaning of a statement is the method of empirically confirming or infirming it" ("Two dogmas," p. 37).⁷ Conventional wisdom is that, if you put verificationism together with the Q/D thesis, you get semantic holism and that *that* is the argument for semantic holism that Quine intends in "Two dogmas."⁸ Whether or not this is the intended argument, it's fallacious. Or so we claim.

Here, to a first approximation, is the way the verificationist argument for semantic holism is supposed to go. The Q/D thesis says that confirmation is holistic; that is, that every statement in a theory (partially) determines the level of confirmation of every other statement in the theory.⁹ Verificationism says that the meaning of a statement is determined by its confirmation relations. The invited holistic inference is that every statement

in a theory partially determines the meaning of every other statement.

Now the first thing to say about this first approximation is that it's unclear – not just in our reconstruction, but in the literature at large – precisely what the force of the conclusion is supposed to be. The metaphor that often goes with semantic holism is as follows: A theory is a sort of network, in which the statements are the nodes and the semantically salient relations among the statements are the paths. The meaning of a statement is its position in the network and is hence defined with respect to the *totality* of the nodes and paths (since the identity condition for networks is itself holistic). If you take this metaphor dead seriously, then any change in the theory changes the semantic value of all the statements that the theory entails; and, strictly speaking, only identical theories can have any of their entailments in common.

That's a reading of the conclusion of the "Two dogmas" argument; but there are reasons to think that it can't be a doctrine that Quine intends. For one thing, as we shall presently have reason to emphasize, among Quine's preferred formulations of the Q/D thesis is that you can hold onto any statement, come what may, in the face of recalcitrant data. But this claim would seem to be literally unintelligible barring a notion of trans-theoretic statement identity of precisely the sort that semantic holism appears to preclude.¹⁰

Second, consider the role of observation statements in confirmation. On the one hand, they must be trans-theoretically identifiable if the *public* character of scientific confirmation is to be preserved. If observation statements are to *decide among* theories, then what your observation statements mean must not depend on which theory you hold. But, on the other hand, it's a truism that if they don't enter into confirmation relations with other statements, observation statements won't confirm anything except themselves. But the Q/D thesis entails that if they do enter into confirmation relations with *any* other statement in a theory, they enter into confirmation relations with *every* other

statement in the theory. And Peirce's thesis says that confirmation relations are constitutive of semantic identity. It presumably follows that if theories differ in any of their entailments, they can't agree (or disagree) about which observation statements are true.¹¹ See re Q's evolution

Formulating a reasonable version of semantic holism (one that is compatible with the Q/D thesis, to say nothing of being compatible with Realism about semantic properties) is, in fact, no small problem. However, it's not *our* problem. For our purposes, we're content to leave the claim that "Every statement in a theory partially determines the meaning of every other statement" largely unexplicated, with the proviso that it be so construed as to entail semantic holism in at least one of the senses we specified in chapter 1: either that the metaphysically necessary conditions for content preclude punctate theories (minds/languages) or that only identical languages can contain intertranslatable statements, or both.

So now we have an argument on the table which infers semantic holism from verificationism together with the Q/D thesis. And the question arises as to what to do about it. Various options suggest themselves. For example, if you are persuaded that the argument is valid, but you don't like semantic holism, it would be reasonable for you to say, "Well then, so much the worse for verificationism (or for the Q/D thesis or for both verificationism and the Q/D thesis)."¹² Given the tribulations to which verificationism, in particular, has been subject over the years, this would not be a wildly unreasonable reaction.

However, though we claim to be as anti-verificationist as the next guy, this is not the course that we propose to follow. Instead, we want to argue that, appearances to the contrary notwithstanding, *even* the conjunction of confirmation holism and verificationism is compatible with the denial of semantic holism (in either of the senses mentioned above). (So this might be a good time to remind the reader that it's not, in our view, entirely obvious that this argument really does occur in "Two dogmas.")

Let's construct an argument, in the spirit of "Two dogmas," to show that a theory (language/belief system) couldn't contain *only* the statement that it's raining:

Premise 1: The statement that it's raining (R) is partially confirmed by the statement that the streets are wet (S). (Metereological platitude.)

Premise 2: Confirmation relations are ipso facto semantic. (Peirce's thesis.)

Premise 3: Statements are individuated by their semantic properties; or, as we will sometimes say, they have their semantic properties essentially. (Truism.)

Lemma: R is individuated, inter alia, by its relation to S.

Conclusion: Any theory that contains R must contain S. A fortiori, no theory could contain just R.

Since this form of argument is clearly independent of the particular choice of examples, and since the Q/D thesis guarantees, at a minimum, that each statement in a theory bears on the confirmation of many, many others, semantic holism follows.

What, if anything, is wrong with this argument? We want to claim that preserving such appearance of validity as it may have depends on being very careful *not to say what a statement is* (and, in fact, in "Two dogmas" Quine is very careful not to say what a statement is). So far as we can see, there are only the following three possibilities, none of which does what the argument needs: Statements are *formulas* or they are *propositions* or they are *formulas together with their conditions of semantic evaluation*. We propose to consider these options in turn.¹³

(1) *Statements are formulas.*

That is, they are morpho-syntactically individuated expressions which have both their semantic properties and their linguistic affiliations *contingently*.¹⁴ It follows that one and the same statement can have many different meanings ("The duck is ready to eat") and that one and the same statement can occur in many different languages ([Empedikli:s li:pt]). If this is the intended reading of "statement," then, for example, premise 1 of the holism argument is equivalent to: "It's raining" is confirmed by "The streets are wet."

Problems:

1.1 This can't be the intended reading; in fact, it's inconsistent with the truism that premise 3 expresses. Formulas precisely do *not* have their semantic properties essentially.

1.2 Formulas as such don't enter into confirmation relations (or entailment relations and the like). It makes no sense to ask whether the form of words "The streets are wet" tends to confirm the form of words "It's raining," since that depends entirely on what these forms of words *mean*.¹⁵ In a language in which "It's raining" means that it's raining and "The streets are wet" means that the streets are wet, the answer is "Yes"; presumably, if a token of the first were true, then, likely enough, a corresponding token of the second would be too. But in a language in which the one means that Chicago is in Indiana and the other means that the cat is on the mat, the answer is presumably "No" (though no doubt it's an empirical issue).

1.3 To identify statements with formulas contradicts (or trivializes) the Q/D thesis under at least one of its preferred formulations: You can hold onto any statement, if confronted by recalcitrant data, by making compensatory adjustments elsewhere in your theory.

That this is a preferred formulation of the Q/D thesis is no accident. Duhem's holism flows in large part from his analysis

of the logic of experimental confirmation. It is not enough, in order for a scientific theory to have testable consequences, that it entail an experimental hypothesis. It must also warrant a host of "auxiliary" hypotheses which will be assumed by the experimental design, auxiliary hypotheses about the operation and the acuity of the experimental apparatus, for one kind of example. The confirmation holism comes in with the recognition that, in principle at least, recalcitrant data can always be accommodated by retaining the experimental hypothesis and abandoning some or all of the auxiliary apparatus. In short, it is largely because auxiliary assumptions are required to generate experimental predictions from experimental hypotheses that no "statement, taken in isolation from its fellows, can admit of confirmation or infirmation at all" ("Two dogmas," p. 41). And it is largely because, in principle, these auxiliary assumptions are always at risk that "our statements about the external world face the tribunal of experience not individually but only as a corporate body" ("Two dogmas," p. 41). (It is precisely at this point in "Two dogmas" that Quine acknowledges his debt to Duhem. See his n. 17.)

But if this is the argument for the Q/D thesis that "Two dogmas" intends, then the formulation that emphasizes the scientist's options about which hypotheses to forfeit is close to the heart of Quine's confirmation holism. And our present point is that it's hard even to make sense of this formulation of the Q/D thesis on a reading that takes "Two dogmas" to mean "formula" when it says "statement."

Here's a slightly different way of putting the issue. The Duhem analysis of experimental confirmation turns upon the thought that one might *give up some hypotheses* (the auxiliary ones) while *holding onto others* (the experimental ones). But, as we remarked a few pages back, this formulation appears to presuppose exactly what semantic holism denies: that there can be a principled, trans-theoretic way of individuating statements, hypotheses, and the like. At first thought, one might seek to reconcile semantic holism with the Q/D thesis by supposing that

when the Q/D thesis says, "You can save any hypothesis if you're willing to pay the price," what it means by "hypothesis" is "formula." Since formulas don't have their meanings essentially, semantic holism allows that their individuation can be principled and trans-theoretic; in particular, that they can be individuated morpho-syntactically.

But second thoughts suggest that this won't do; when the Duhem analysis talks of "hypotheses," it can't be just forms of words that it has in mind. For, consider:

1.3.1 What is it for a *form of words* to be confronted by data (recalcitrant or otherwise)?

1.3.2 It's only epistemologically interesting that you could hold onto "Burning is the liberation of phlogiston" even in the face of Lavoisier's results if "Burning is the liberation of phlogiston" means that burning is the liberation of phlogiston. It's no news that you could hold onto it in the face of those results if it means that Greycat has whiskers.

1.3.3 If statements are just formulas, you don't have to "make compensatory adjustments elsewhere" in order to hold onto them; if a formula that you like gets into trouble, use it to mean that two and two is four and leave the rest of the theory alone.

We conclude that reading "statement" as "formula" precludes a substantive reading of the Q/D thesis.

1.4 If statements are formulas, then translation holism looks to be trivially false. For, clearly, otherwise different languages *can* contain the same *formulas*: English and German both contain [Empedikli:s li:pt]. And, if it is possible for there to be a language which contains only one *morpho-syntactically* well-formed form of words, then content holism is false too — even if, in such circumstances, the form of words wouldn't *mean* anything.¹⁶

We conclude that statements can't be formulas.

(2) *Statements are propositions.*

According to this account, statements are trans-theoretical (trans-linguistic) entities, entities that can be expressed by the formulas of more than one theory. (In particular, they're the sorts of things that the "that" clauses in "the fact that P" or "the belief that P" are said to name.) So questions like "How do you spell the statement that P?" or "What's the first word in the statement that P?" are nonsensical, though the question "How do you state that it's raining in German?" is perfectly OK.

This would perhaps be the natural way to read "statement" if the present discussion were about confirmation rather than meaning, since it's natural to say things like "That the streets are wet confirms (suggests/shows/makes it seem likely, and so on) that it's raining." Moreover, this reading of "statement" secures the truism that statements have their conditions of semantic evaluation necessarily. It's a nonaccidental property of the proposition that it's raining that it's true iff it's raining.

Problems:

2.1 Though, as we remarked, we aren't primarily interested in *ad hominem*s, it's worth remarking that this clearly can't be what Quine means by "statement." It's precisely because Quine is suspicious of language-independent meanings, propositions, and the like that he's suspicious of the *a/s* distinction. Indeed, suspicion of one is really just suspicion of the other. So the idea that statements are propositions, in the sense of language-independent meanings, flies in the face of the first part of "Two dogmas."

2.2 Translation holism is a constraint on the expressive power of *languages, theories*, and the like (roughly, of systems of symbols). It says, for example, that if L can express one of the statements that English can express, then L can express practically all the statements that English can express. But if statements are propositions, the putative Quinean argument for holism doesn't have anything like this consequence. The

argument does perhaps show that there are internal relations among *propositions*; for example, if there is the proposition that it's raining, then, according to the argument, there must be the proposition that the streets are wet. But from the fact that there being the proposition that P necessitates there being the proposition that Q, it does not follow (at least, it doesn't follow without further argument) that a language can't *express* the proposition that P unless it can also express the proposition that Q.

This is a perfectly general point, worth stressing outside the present context. Here's a bad argument for content holism about concepts:

The concept *cat* is partially constituted by its connection to the concept *animal*; nothing could *be* the concept *cat* unless its applies only when the concept *animal* applies. So no mind entertains the concept *cat* unless it entertains the concept *animal*. So a mind that entertains the concept *cat* can't be punctate.

This argument is a non sequitur if you read "concepts" the way philosophers usually do: namely, as abstract objects which individuals can *share*.¹⁷ Reading "concepts" that way, there is no obvious inference from "There wouldn't be the concept A but that there is the concept B" to "Entertaining the concept A requires entertaining (being able to entertain, being able to learn to entertain, grasping, and so forth) the concept B." In fact, barring further argument, assumptions about necessary relations among concepts don't appear to have any *psychological* consequences at all;¹⁸ nor, for precisely parallel reasons, would they appear to constrain the expressive power of languages in which the concept can be articulated.

2.3 We take this to be the really crucial consideration: If statements are propositions, then statements have their contents essentially; propositions are individuated by reference to their contents. So, if contents are means of confirmation (as per

Peirce's thesis), then statements have their means of confirmation essentially. But the Q/D thesis is (inter alia) the idea that statements have their confirmation conditions *contingently*; according to the Q/D thesis, it's *contingent* what confirms what. So you can't consistently hold both Peirce's thesis and the Q/D thesis if you also hold that statements are propositions. (More of this sort of argument in the next section; it shows that statements can't be formulas together with their conditions of semantic evaluability either.) So Quine couldn't hold that statements are propositions even if he were a friend of propositions – which, of course, he's not.¹⁹

So statements can't be propositions.

(3) *Statements are formulas together with their associated conditions of semantic evaluability.*²⁰

This seems to us to be the most interesting candidate. In particular, it allows a substantive construal of the Q/D thesis, which now says that you can hold onto a form of words, come what may, *without equivocation*; and this is an interesting claim, one that isn't at all self-evident.

So then, what do you get if you conjoin the Q/D thesis with Peirce's thesis on this construal of "statement?" Presumably something like this: Statements have their conditions of semantic evaluation essentially; Peirce's thesis says that what it is for a statement to have the conditions of semantic evaluation that it does is for it to have the confirmation relations to other statements that it does (at a minimum, conditions of semantic evaluation supervene on confirmation relations); so statements have their confirmation relations to one another essentially.

Much of this may seem to be in the spirit of "Two dogmas," but it can't be what Quine (or anyone else who accepts the Q/D thesis) has in mind either.

3.1 A formulation of the Q/D thesis that is very close to its core is this: Confirmation relations are a posteriori; it's *a matter*

for scientific discovery what confirms what,²¹ and we change our estimates of confirmation relations to accommodate the evidence, just as we do our other theoretical commitments. Indeed, as we've already seen, the heart of Quine's argument against what he calls "reductionism" (the doctrine that statements are analytically connected to their confirmation conditions) is, to all intents and purposes, that it *can't* be a matter of meaning that "The streets are wet" confirms "It's raining," because we might simply *discover* that it doesn't. But if confirmation relations among statements are revisable and if meaning is construed in terms of confirmation, then statements don't have their semantic properties essentially. So statements can't be formulas *together with their conditions of semantic evaluation*. (This is the same argument that showed that statements can't be propositions; see above.)

The reader might well be wondering whether there is *any* version of Peirce's thesis that can be squared with the idea that confirmation relations are revisable. After all, Peirce's thesis just *is* the claim that confirmation relations constitute semantic relations and are therefore not contingent. Quite so. The arguments we have been considering amount to a dilemma:

If statements as such are just formulas, you can't make sense of talk about confirmation relations holding among them.

If statements as such are semantically interpreted, then the Q/D thesis says that they have their confirmation relations contingently and Peirce's thesis says that they have their confirmation relations essentially.

It looks as though what the Q/D thesis and Peirce's thesis say about statements are mutually inconsistent. We conclude that, questions about semantic holism to one side, there must be something deeply wrong with "Two dogmas," since it is explicitly committed to both principles.

It is, *prima facie*, inadvisable to try to run an argument for

semantic holism (or, indeed, for anything else) which takes both the Q/D thesis and Peirce's thesis as premises. You could try the following, however. The Q/D thesis requires confirmation conditions to be *a posteriori*; but, strictly speaking, it does not require them to be *contingent*. It is thus compatible with the Q/D thesis that confirmation relations should be (not linguistically but) *metaphysically* necessary. On the one hand, *if P confirms Q, then it's necessary that P confirms Q*; but, on the other hand, since metaphysical necessities are presumably known *a posteriori*, only inquiry would allow us to determine when a confirmation relation holds.

Given his empiricism (to say nothing of his anti-essentialism), it isn't easy to imagine Quine opting for this markedly unepistemological reading of "confirmation condition"; but at least it does square the Q/D thesis with Peirce's thesis. Notice, however, that it doesn't help with the argument for semantic holism. On the contrary, semantic holism *clearly* would not follow from confirmation holism according to this account. For, though what a statement means now depends on what confirms it, what confirms it does *not* depend on the theory it's embedded in. On the present assumptions, it could perfectly well turn out that the embedding theory for P might be *entirely wrong* about what the confirmation conditions of P are, in which case P's role in the theory (including what the theory says about P's confirmation conditions) would be simply *irrelevant* to what P means. The moral is that to get semantic holism from confirmation holism, you need to relativize what a statement means in a theory to what *that theory* says about the confirmation conditions of the statement.²²

3.2 A second, relatively minor reason not to suppose that what Quine means by "statement" is *formula with its conditions of semantic evaluation* is that this interpretation would fail to secure the *immanence* of confirmation. (See *Philosophy of Logic*, pp. 19–20; we are extending Quine's usage from languages to theories.) Immanence is the idea that because confirmation is defined over sorts of entities whose connection

to a particular theory (/language) is essential, it need not be possible to construe such questions as whether two theories agree about confirmation relations. This is, pretty clearly, a thesis that Quine wishes to maintain. But immanence fails on the present proposal, just as it does on the proposal that statements are formulas. There is no obvious reason why the formula "It's raining" couldn't, by accident, turn out to be a Swahili sentence that is true iff it's raining; in which case, English and Swahili would have a statement in common.

So statements can't be formulas together with their conditions of semantic evaluation. So, apparently, there is nothing that statements *can* be, consonant with the use to which the "Two dogmas" argument for semantic holism wants to put them. So the argument is unsound.

What's gone wrong? We think it's this: The strategy we've been attributing to "Two dogmas" is to infer semantic holism from confirmation holism. In order to do so, it must take for granted that the X's that enter into confirmation relations (in particular, the X's to which the Q/D thesis applies) are the very same X's that semantic theories are about; they're the very same things whose conditions of semantic evaluation semantic theories specify.²³ The trouble is, however, that whereas the natural objects of semantic interpretation are linguistic entities like formulas, the natural bearers of confirmation relations are trans-linguistic entities like propositions. So, even though confirmation holism is quite likely true, and even though verificationism is assumed for the sake of argument, there is no sound inference from those premises to semantic holism, because confirmation holism and verificationism are true of *different things*.²⁴

We conclude this section by reminding the reader that, since we don't think that verificationism *is* true, the previous argument is a little on the hypothetical side. But we sense, in the philosophical community at large, some sympathy for the idea that, well, maybe verificationism is sort of *a little* true;²⁵ at least, it might be true enough to buy you a semantics from

which meaning holism can be inferred. So it seemed to us worth stressing that the (putative) Quinean argument has troubles that run deeper than the odd false premise. It is – ineliminably, in our view – a fallacy of equivocation.

ANOTHER ARGUMENT FOR MEANING HOLISM

We're almost done with the discussion of meaning holism in "Two dogmas." But before we pack up the chapter, it's worth mentioning a curious subsidiary argument that Quine suggests both in "Two dogmas" and in "Epistemology naturalized" (see p. 72): namely, that "Russell's concept of incomplete symbols defined in use" implies *statement* holism; that is, it implies that "the primary vehicle of meaning . . . [is not] . . . the term but in the statement" ("Two dogmas," p. 39). This suggestion is worth considering in the present context, because a demonstration that the unit of meaning is no smaller than a whole statement (or sentence) might be a lemma on the way to demonstrating that the unit of meaning is no less than a whole theory (or language). Also, it's possible to wonder how the statement-as-semantic-unit story is to be squared with the widely prevalent idea that sentences (and, hence, the statements that they are used to make) have *compositional* semantic structure; that is, that the meanings of sentences are derived from the meanings of their lexical constituents.

Unfortunately, Quine doesn't say how the argument from definitions-in-use to statement holism is supposed to run; and it's not easy to guess what he could have had in mind. For one thing, statement holism is presumably a *modal* claim, something like that no expression *could have a semantic value* except as it functions as part of a sentence. The least Quine would appear to need to get this modal conclusion would be the corresponding modal premise: namely, not just that (some) terms *are* defined in use but that all terms *have to be*. It is, to put it mildly, unclear

that Russell's remarks on definite descriptions warrant any such claim.

In fact, it's not clear whether the (presumed) facts about definitions-in-use warrant *any* claim about *meanings*. Let's suppose that some words are defined with respect to their sentential contexts, just as Russell thought. It remains open – and crucial – *which aspects* of their sentential contexts these words are defined with respect to. In particular, it depends on whether words that are defined in use are ipso facto defined relative to *semantic* properties of their contexts.

Consider, for example, [kri:ks], which means one thing in "The door . . ." but something quite different in "The . . . flow." Do these considerations show that [kri:ks] is not a "unit of meaning" (that is, that the unit of meaning is not [kri:ks] but, as it might be, [kri:ks]-in-a-sentence-about-doors)? Well, perhaps; but also perhaps not. Maybe the right story is that the units of meaning are [kri:ks]_{verb} in the one case and [kri:ks]_{noun} in the other. If this *is* the right story, then it suggests that the sentence is the unit of *syntax*; that is, that words couldn't have the *syntactic* properties they do if they did not occur as constituents of sentences.)

This suggestion seems plausible enough; on the face of it, syntactic properties look to be ones that words have in virtue of their relations to the sentences that contain them.²⁶ By contrast, it is *by no means* obvious that their *semantic* properties are ones that words have in virtue of their relations to sentences that contain them. And it is also not obvious that the questions "What's the unit of semantic analysis?" and "What's the unit of syntactic analysis?" have to get the same answer.²⁷)

FACTS ABOUT MEANING?

So much for the arguments for semantic holism in "Two dogmas"; we can now say what the rest of this book will be about. Since it appears that the arguments usually attributed to

"Two dogmas" are deeply flawed, the question arises as to whether there are other arguments that are more convincing. In the course of discussing this question, it will often turn out that whether Quine is right about the *a/s* distinction is crucial to the direction that the polemics take. We propose, if only because it is the more conservative policy, not to appeal to the *a/s* distinction in the evaluation of holistic theses about meaning; in effect, we will assume that Quine is right and that the *a/s* distinction cannot be sustained.²⁸ So then, if there is something wrong with semantic holism, it is *not* that only the statements that *S* is *analytically* related to are constitutive of what *S* means.

We want to emphasize, however, that conceding this is conceding very much less than that (for example) there are no facts about meaning or that there are no truths that hold "in virtue of meaning alone" or that the notion of a semantic rule is unintelligible or (a fortiori) that there can be no such thing as a semantic theory. It's perfectly possible to be eliminativist about an *epistemologically based a/s* distinction but not about semantics as such and, indeed, not about the *a/s* distinction as such.

To see what's at issue, imagine somebody who holds a reductionist theory of meaning; for example, somebody like Skinner, who holds that for "dog" to mean *dog* in a certain speaker's mouth is for the speaker to have the habit of uttering "dog" when there are dogs around. Perhaps it needn't be said that we think this sort of theory is unlikely to work. But there are *some* lines of argument against which it's defensible; and here's one. Look, someone might say, if Skinner were right, then you could have a situation in which a speaker has two responses ("dog" and "shmog," as it might be) that are conditioned to *exactly the same stimuli*. But then it would follow that these responses would be *synonymous* for that speaker. So then the following sentence would be *analytic* in the speaker's language (assuming he has the logico-syntactic apparatus to frame it): "Whatever is a dog is a shmog." But Quine showed in "Two dogmas" that there are no such things as synonyms or analytic

sentences. So Skinner's semantics *must* be wrong. A priori! In fact, *all* semantic theories must be wrong, a priori, except for the nihilistic theory which says that there are no semantic properties.²⁹

What has gone wrong this time? We take it that, strictly speaking, Quine in "Two dogmas" did *not* show, or even argue, that there are *no* semantic facts, or even that there are no analytic truths.³⁰ Rather, what we are conceding is that if there *is* sense to be made of meaning and the associated notions, it can't be reconstructed by reference to *statements that a speaker assents to*. Or, equivalently for these purposes, if Quine is right in "Two dogmas," then what you mean can't be reduced to what inferences you are prepared to accept.

For example, it can't be that whether you mean that John is a bachelor by your utterance of "John is a bachelor" depends on whether, if you are prepared to utter "John is a bachelor," then you are prepared to utter "John is an unmarried man," or on whether you are prepared to accept the inference that if "is a bachelor" applies to John, then "is an unmarried man" does too. Because, Quine argues, what inferences you are prepared to accept (and/or what you are prepared to utter given that you are prepared to utter . . .) depends not only on what you intend your words to mean, but also on how you take the (non-linguistic) world to be. And there is no principled way to separate the respective contributions of these factors. Knowing which inferences someone accepts doesn't tell you which of them he accepts a priori; so it doesn't tell you which of them is analytic.

Quine's rejection of analyticity, insofar as it's actually argued for (and insofar as we are proposing to concede it), is a rejection only of the possibility of an *epistemic* criterion for "true in virtue of meaning." In principle, at least, everything else remains wide open. For example, it's left open that you might be able to reduce semantic relations to resemblance relations (the way Hume wanted to) or to conditioning relations (the way Skinner wanted to) or to nomological relations (the way

Dretske wants to) or to nonepistemic intentional relations of being appeared to (the way we gather phenomenologists want to) – or to other relations as yet undreamed of. Any of these reductions would imply corresponding notions of synonymy/analyticity (Ideas that resemble the same things are the same Ideas; words that are conditioned to the same things are synonyms, and so forth). And, since none of these notions of synonymy/analyticity is epistemic, none of them is in jeopardy of the sorts of considerations that Quine offers in “Two dogmas.”

On this view, the only doctrine that *can't* be held, consonant with Quine's rejection of the *a/s* distinction, is the following conjunctive one: Some inferential relations are constitutive of semantic relations, *and* which they are can be determined by applying an epistemic criterion like aprioricity or unrevisability. Inter alia, meaning holism remains open; and so too does a reductionistic atomism (and so too, of course, does semantic nihilism).

What we now want to know is: What can someone who *isn't* a verificationist say on behalf of the first of these options?

3

DONALD DAVIDSON: MEANING HOLISM AND RADICAL INTERPRETATION

The most interesting of the arguments for holism to be discussed in this chapter and the following two pursue a common strategy. The idea is to show that certain principles of charity are constitutive of content ascription. It's purported that these principles are intrinsically holistic, so the holism of the intentional is entailed.¹

All the arguments we'll be looking at are more or less transcendental in style; but it is characteristic of the ones to be considered in this chapter that they have premises that are simultaneously epistemological and metaphysical. The basic idea is that it is metaphysically constitutive of facts about content that they must be accessible to someone in the epistemological situation of a *radical interpreter* and that radical interpretation is impossible unless principles of charity are invoked. Our primary sources for this discussion will be Donald Davidson's "Truth and meaning," "Reply to Foster," and "Radical interpretation."

MEANING THEORIES, TRUTH THEORIES, AND RADICAL INTERPRETATION

Philosophers have traditionally disagreed not only about what the right theory of meaning is for a natural language (for

NOTES

NOTES TO CHAPTER 1 (pp. 1–6)

1. By which we intend “not atomistic” rather than “of or pertaining to anatomy.”
2. We will, throughout, use the authorial “I” and the authorial “we” interchangeably, as ease of exposition dictates. “He” and “his” are often used without implication of gender.
3. We’ll generally save “meaning holism” or “semantic holism” for the broader, and less precise, doctrine that meaning is somehow holistic. So meaning holism is true if (but perhaps not only if) either content holism or translation holism is true. The main reason for bothering to distinguish content holism from translation holism is that a meaning holist might admit the possibility of punctate languages, minds, and the like as a sort of metaphysical curiosity but still deny that a punctate language could express anything that can be expressed in English, the idea being that in *nonpunctate* languages (like English) the meanings of sentences are constituted by their relations to one another. Content holism thus precludes possibilities that translation holism leaves open. We will call the doctrine that asserts translation holism but allows punctate languages “semi-holism” when it’s important to distinguish it from other holistic options.

For many purposes, however, the various kinds of meaning holism tend to stand or fall on much the same considerations; we will, therefore, often run them together.

4. This way of putting things depends on allowing the notion of a language itself to be construed relatively narrowly – as a set of sentences, say, rather than a life-style. Anthropological holism is, in part, the idea that

this narrow reading of “language” is hopelessly artificial and that, in the long run, there is no real distinction between what is linguistic behavior and what isn’t or, ultimately, between languages and whole cultures. That may be right. If it is, then a lot of linguists have been wasting their time barking up phrase-structure trees. We remark in passing, however, that it is possible to imagine a view that is holist in the broad, anthropological sense but nevertheless leaves open the possibility of punctate languages. For example: symbols get their meanings from the way they are embedded in Forms of Life, but there’s no internal connection between being so embedded and being part of a symbol *system* (for example, being part of a language with a compositional syntax and semantics). The “primitive languages” that Wittgenstein imagines in the early paragraphs of the *Philosophical Investigations* are, perhaps, meant to be holistic in the broad but not the narrow sense.

5. Where Frege himself stands is a little unclear. On the one hand, it’s a famous Fregean view that words have meaning only as constituents of (hence, presumably, only in virtue of their use in) sentences, and this view looks to be inherently anatomic; but, on the other hand, Frege certainly thought that the semantics of sentences is compositionally determined by the semantics of the words they contain (plus their syntax), and this suggests that lexical semantics must in some sense be prior to sentence semantics. Whether, and in exactly what way, these doctrines can be reconciled is a notorious crux in Frege interpretation.
6. This sort of issue isn’t made to go away by taking the objects of theoretical interest to be idiolects rather than languages (as, indeed, many linguists are inclined to do, even at the cost of denying that the basic function of natural languages is to mediate communication between its speakers; see Chomsky, Halle, and others). For there is still the problem of communication between *time slices of an idiolect*; if holism is true and idiolects are the minimal units of meaning, how could I have incrementally learned the idiolect that I now speak?

Davidson suggests that “we cannot accurately describe the first steps towards the conquest [of a language] as learning part of the language; rather it is a matter of partly learning” (“Theories of meaning and learnable language,” p. 7). That is, it is possible for a child to *partially learn a language* without learning *part of the language*. This is not, however, a suggestion we claim to fully understand.

7. We only say it *might* turn out this way. We’re currently running the discussion on the assumption that there is an argument from the premise that semantic properties are *anatomic* to the conclusion that they are *holistic*. But precisely which holistic consequence follows from the assumption that a semantic property is anatomic depends, of course, on

exactly how this argument is supposed to go. We will return to this question presently.

8. The argument that properties like R^* are holistic often assumes that *meaning* is holistic and that meaning determines reference. Whether reference holism can be defended without this assumption is a question of great philosophical interest, but not one that we will consider in this book.
9. It wouldn’t follow from R^* ’s being holistic that theories are incommensurable unless their ontologies are *identical*. There might be some sense of “similar” in which theories are commensurable if their ontologies are similar enough. (We’ll discuss this sort of possibility presently.) Nor would it follow that if T_2 and T_3 are both commensurable with T_1 , then there are things that T_2 and T_3 can both refer to. The ontological requirements for commensurability might permit that T_2 is commensurable with T_1 because they can both refer to a’s, b’s, and c’s and T_3 is commensurable with T_1 because they can both refer to d’s, e’s, and f’s; that is, having expressions that refer to a, b, and c and having expressions that refer to d, e, and f are both sufficient for sharing the ontology of T_1 , though neither is necessary. (This sort of possibility was pointed out to us, in a slightly different context, by Paul Boghossian, Barry Loewer, and Tim Maudlin; see below.)
In either case, the urgent issue for Scientific Realism is whether there is, short of identity, a *principled* answer to the question “Which sorts of overlaps between ontologies are sufficient for empirical commensurability?” In the terminology of Kuhn’s *The Structure of Scientific Revolutions*, this is approximately the question as to whether it’s principled when different theories belong to the same paradigm.
10. Paul Churchland appears to hold, largely on the ground that properties like R^* are holistic, that only the final, literally true physics will be able to refer to *anything at all*. Churchland seems to take this view to be good news for Realism, but it’s not clear to us why he does. On his account, the only science that has an ontology – a fortiori, the only science for whose ontology reality can be claimed – is not one that any human scientist is ever likely to profess.
11. We make a point of *not* using shopworn examples like “If you believe ($P \rightarrow Q$) and P , then you believe Q .” These sorts of generalization presuppose a notion of *identity and difference* of belief content rather than a notion of belief content *per se*.
12. This is a bullet that is frequently bitten. For example, Field holds that the meaning of a sentence is determined by its “referential meaning” together with its “conceptual role” (“Logic, meaning and conceptual role,” p. 390). As he recognizes, the conceptual role part implies semantic holism. (Field

shares the usual doubts about the a/s distinction.) Field accepts the consequence that his semantics is therefore

compatible with a great deal of pessimism about the clarity of the notion of inter-speaker synonymy. . . . My own inclination is not to try to provide such an account but to learn to live without the concept of inter-speaker synonymy, and all other concepts in terms of which inter-speaker synonymy could be defined. (The place that such concepts appear to be needed is in belief–desire psychology. I believe that any such psychology formulated in terms of such concepts can be reformulated so as not to employ them and that there are independent grounds for preferring the reformulated theory. (Ibid., pp. 398–9)

Field doesn't, however, say how this reformulation is to be achieved.

In a quite different context, but a rather similar spirit, Roy Harris remarks that "It is arguable that if translation is taken as demanding linguistic equivalence between texts, then the Saussurean [structuralist] position must be that translation is impossible" (in Saussure, *Course in General Linguistics*, p. xiii). Harris does not take this to be a *reductio ad absurdum* of the Saussurean position.

13. A generalization is "robust" to the extent that the individuals that fall under it are otherwise heterogeneous in lots of ways; correspondingly, a definition is robust if it is satisfied under lots of otherwise heterogeneous conditions, and so forth.
14. This is the received account of Davidson's view, but Davidson's view may be more nuanced than the received account supposes. Davidson clearly holds that there can't be *exceptionless*, or "*homonomic*," intentional laws; but it wouldn't *seem* to follow from this that intentional laws can't support counterfactuals, back singular causal truths, and so on. After all, the (presumably) heteronomic character of geological laws doesn't prevent them from doing so. What Davidson takes the bottom line on these topics to be is not something we're at all sure about. See Lycan, "Psychological laws"; Rosenberg, "Davidson's unintended attack on psychology"; Dennett, "Mid-term examination: compare and contrast."
15. It's sometimes pushed pretty hard that the holistic and systematic character of the semantic (/intentional) isolates hermeneutical investigations from modes of criticism that are pertinent elsewhere. Sometimes it's pushed to the verge of mysticism:

The logic of difference is a non-self-identical logic, one that eludes all the normative constraints which govern classical reason. If language

is marked by the absence of "positive terms" – if meaning is differential through and through – then any theory which attempts to conceptualize language will find itself up against this ultimate limit to its own explanatory powers. (Norris, *Derrida*, p. 91)

16. For example, here's Gilbert Harman in *Thought*:

Two people can be said to mean exactly the same thing by their words if [sic; "only if"?] the identity-translation works perfectly to preserve dispositions to accept sentences under analysis and actual usage. To the extent that the identity-translation does not work perfectly, people do not mean *exactly* the same thing by their words; but if the identity-translation is better than alternatives we will say that they mean the same thing by their words. Here we mean by *the same*, *roughly the same* rather than *exactly the same*. . . . The only sort of sameness of meaning we know is similarity in meaning, not exact sameness of meaning. This is where the defender of the analytic-synthetic distinction has gone wrong; he confuses a similarity relation with an equivalence relationship. (pp. 109–10)

17. Twin worries (à la Putnam, "The meaning of 'meaning'") are not the issue here; choose any physical state of affairs, relational or otherwise, on which you are prepared to believe that belief systems supervene.
18. Notice that this is much the same problem as has led so many philosophers to despair of the project of constructing a robust notion of content *identity* by appealing to some suitably abstract notion of identity of inferential role. Some inferences (traditionally the analytic ones) count, and some inferences (traditionally the synthetic ones) don't count, and there appears to be no principled way of saying which are which. This problem *does not disappear* if you replace "count"/"don't count" with "count much"/"don't count much."
19. Nor, of course, would the more sanguine conclusions that are often drawn from meaning holism – as, for example, that the assumptions of commonsense Intentional Realism are immune to challenge from the physical sciences. See above.
20. Others deny the a/s distinction, accept that holistic consequences are entailed, and argue that the right moral to draw is that there really aren't any intentional properties. Quine takes this line in certain of his moods, and so do Dennett and Stich in certain of theirs; the Churchlands take it all the time.
21. This is one reason for being skeptical as to whether Quine's "Two dogmas of empiricism" contains an argument for semantic holism along

- the lines of A, though it is widely interpreted as doing so. See the next chapter.
22. The relevant consideration is this: If A is a proposition that you have to believe to believe P, then presumably $P \rightarrow A$ must be analytic. (If nobody could believe that something is a dog unless he believed that that thing is an animal, then the belief that if something is a dog, then it's an animal is a semantic truth.) According to the present assumptions, however, there is *no* proposition that you must believe in order to be able to believe P. (Either believing A or believing B is sufficient, but neither believing A nor believing B is necessary.)
 23. More precisely, it closes the book against the possibility of atomism about *belief*. This is a distinction we dwell on in chapter 4, q.v.
 24. It's worth emphasizing, in the current atmosphere of near universal holistic consensus, that until very recently, and for a very long while, the philosophical consensus for semantic atomism seemed equally secure. We commend this historical reflection to philosophers who say that no argument for semantic holism is required because it is self-evidently true, or that the atomism of semantic properties is intuitively obvious.
 25. The variations on this theme in the secondary literature on semiotics are endless. Here are examples, chosen practically at random:

It is a cardinal precept of modern (structural) linguistics that signs don't have meaning in and of themselves, but by virtue of their occupying a distinctive place within the systematic network of contrasts and differences which make up any given language. (Norris, *Derrida*, p. 15)

For it is a major precept of modern structural linguistics that meaning is not a relation of identity (sic!) between signifier and signified but a relation of differences, the signifying contrasts and relationships that exist at every level of language. (Ibid., p. 85)

The choice is thus between a linguistic atomism that grounds meaning in a language/world ("sign"/"signifier") relation (though not, one might have thought, an *identity* relation) and a linguistic holism which grounds meaning in the relation between a symbol and its role in a language; and "modern structural linguistics" teaches us to prefer the second option.

This is, of course, a wildly tendentious account of what linguistics teaches us about meaning. Consider how badly it comports with model theoretic, or situational, approaches to the semantics of natural languages, all of which assume that language/world relations (like "satisfaction," "extension," and "denotation") are what the theory of meaning is about.

- (In particular, they assume that it's about how the syntactically complex expressions in a language inherit these language/world relations from their syntactically simpler constituents.)
26. Philosophical interest in resemblance theories of meaning much pre-dates the British empiricists, of course. See Plato, *Cratylus*, and Aristotle, *De Interpretatione*.
 27. For a brief discussion of why resemblance theories don't work, see Fodor, *The Language of Thought*, ch. 4. For the classic discussion of why conditioning theories don't work, see Chomsky, "Review of B. F. Skinner's *Verbal Behavior*."
 28. For a discussion of some recent attempts to construct an atomistic theory of content, see Fodor "A theory of content."

NOTES TO CHAPTER 2

1. We will follow what we understand to be Quine's usage, according to which *reductionism* is a species of *verificationism*. (What precisely the distinction between the two amounts to will be discussed below.) Readers who are accustomed to use "reductionism" to name a type of *ontological* theory should bear in mind that Quine's usage is eccentric.
2. An ontological – specifically, an anti-Realist – construal of the pragmatism in the last pages of "Two dogmas" certainly seems natural. But, on a close reading, it is less than fully apparent that that's what Quine actually intended. For example, though Quine says that gods and physical objects are both just "cultural posits," the explicit claim is only that they are comparable "*epistemologically*" (our emphasis). The ontological moral – if, indeed, there is supposed to be one – is pretty carefully not drawn.
3. Another version of the Q/D thesis says that "the unit of confirmation is the whole theory"; and this *doesn't* follow from these Realist considerations. But we doubt that Quine actually holds the Q/D thesis in this latter form. Glymour remarks that "[even] without analytic truth we need not . . . defy history and good sense by insisting that evidence must bear on all of a theory (let alone on all of science) or none of it or that we must accept or reject our theories as a single piece" (Glymour, *Theory and Evidence*, p. 152). Glymour's point is that, given recalcitrant data, we can pick and choose which bit of theory to give up; we don't have to give it all up. Glymour is surely right about this; but it's far from clear to us that Quine intends to deny it. Quine's claim isn't that if you get recalcitrant data,

Epistemic Assimilation

everything has to go; it's that *what goes and what stays is rationally up for grabs*. Specifically, what goes and what stays can't be decided a priori by appeal to semantics.

4. The reference to Frege is, however, labile. In the original *Philosophical Review* version of "Two dogmas," the reference is to Russell; in the first edition of *From a Logical Point of View*, the reference is to Frege; and in later editions of *From a Logical Point of View*, the reference is to Bentham. For our present purposes, any of the three will do.
5. Since Quine exegesis is always dangerous territory, perhaps we'd best quote some philosophers who read the text in this way. Here is Putnam: "Quine's argument for meaning holism in 'Two dogmas of empiricism' is set out against the meaning theories of the positivists . . . Quine argues that . . . individual sentences are meaningful in the sense of making a systematic contribution to the functioning of the whole language" ("Meaning holism," p. 405). Similar views are expressed in Putnam, "The analytic and the synthetic," to which Quine refers approvingly in *Word and Object*, p. 57.

Consider also the following from Gibson (*The Philosophy of W. V. Quine: An Expository Essay*) (a source which Quine commends for its "full understanding" of his work).

Let's call Quine's version of Peirce's thesis (P):

P: The meaning of a sentence turns purely on what would count as evidence for its truth.

(P) is a form of verificationism. It is the thesis that equates the concepts of meaning and evidence. . . . Let us call Quine's version of Duhem's thesis (D).

D: Single theoretical sentences do not always or usually have a separable fund of *evidence* to call their own; they have their funds of evidence only when connected to larger blocks of theory.

Now, if (P) is regarded as a statement of the equivalence of the concepts of meaning and evidence, then (D) could as well be read as:

(D') Single theoretical sentences do not always or usually have a (separable) *meaning* to call their own; they have their meanings only when connected to larger blocks of theory. (pp. 80–1)

For other philosophers who more or less explicitly hold (or hold that Quine holds, or both) that verificationism together with the Q/D thesis entails semantic holism, see P. S. Churchland, *Neurophilosophy: Toward an Unified Theory of the Mind/Brain*, pp. 265–7; P. M. Churchland,

Matter and Consciousness: A Contemporary Introduction to the Philosophy of Mind, ch. 3; Gilbert Harman, "Meaning and semantics," pp. 11–12.

6. "Quine's theory, as presented at the end of 'Two Dogmas,' is verificationist, since it has entirely to do with the effect which experience has on what we hold to be true, as opposed to the truth of what we say independently of whether we have reason for it" (Dummett, *Frege*, p. 592).

Notice, by the way, that the identification of meaning with means of confirmation needn't be viewed as a metaphysical *reduction* of either to the other. The morning star doesn't *reduce* to the evening star; it just is it.

7. See also Quine, "Epistemology naturalized":

If we recognize with Peirce that the meaning of a sentence turns purely on what would count as evidence for its truth . . . then the indeterminacy of theoretical sentences is the natural conclusion. . . . Should the unwelcomeness of the conclusion persuade us to abandon the verificationist theory of meaning? Certainly not. The sort of meaning that is basic to translation, and to the learning of one's own language, is necessarily empirical meaning and nothing more. . . . Surely one has no choice but to be an empiricist so far as one's theory of linguistic meaning is concerned. (pp. 80–1)

Not only *is* Quine a verificationist; he thinks there aren't any other options.

We're stressing this because many philosophers apparently read "Two dogmas" as *rejecting* verificationism. This is because they fail to notice a distinction that it is centrally important to keep in mind. As we've seen, what is rejected in "Two dogmas" is not *verificationism* but what Quine calls "reductionism": namely, the theory that there is, for each statement, a corresponding range of confirming conditions determinable a priori. Reductionism is inherently localist with regard to confirmation; whereas verificationism, as such, is neutral on whether confirmation is holistic. The former is explicitly attacked in "Two dogmas," but there's every reason to suppose that rejecting verificationism is an option Quine would never have considered.

8. See, for example, Loar:

The confirmation or verification of a sentence is, for Quine, undetachable from its potential evidential connections with virtually any other sentence. Within the positivist tradition, in which meaning and evidential responsiveness are identified, that epistemological

theory naturally leads to a holistic conception of meaning. . . . It is quite in . . . [this] . . . spirit to take each sentence to have a kind of meaning, identified not with its own empirical meaning, but with the totality of its evidential connections to other sentences and perceptual stimuli. (“Conceptual role and truth conditions,” p. 273)

See also Fodor, *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, ch. 3.

9. The point isn't just that you can always imagine a state of affairs such that, *if it obtained*, then the truth of, as it might be, “Grass is green” would bear on the truth of, as it might be, “Mars has rings.” Since *potential* confirmation relations, in this sense, presumably hold between *any* contingent statements, it's unclear how a plausible version of Peirce's thesis could reconstruct semantic relations in terms of them. The point is rather that the level of confirmation of a theory (and hence, derivatively, of each of its entailments) depends in part on such considerations as simplicity, conservatism, and the like, which, because they are global, are simultaneously sensitive to *all* the theory's commitments.
10. Conversely, as Dummett has remarked, Quine likes an ecumenical story according to which the connectives *mean different things* in classical and intuitionistic logic. But, again, barring some trans-theoretic notion of statement identity, it's unclear how you decide *which* connectives it is that the two kinds of logic assign different meanings to. (The thought that there are perhaps *no* theorems that classicists and intuitionists can both accept without equivocation seems, to put it mildly, unintuitive.)
11. Because of this problem, Quine's post-“Two dogmas” writings seem increasingly inclined to treat observation statements as exceptions to the Q/D thesis; that is, as having their confirmation conditions *locally* determined. See, for example: “The observation sentence, situated at the sensory periphery of the body scientific, is the minimal verifiable aggregate; it has an empirical content all its own and wears it on its sleeve” (“Epistemology naturalized,” p. 89). Heaven only knows how this is supposed to square with the version of the Q/D thesis according to which *any* sentence may be rationally abandoned given sufficient pressure from global simplicity, coherence, conservatism, and the like. In any event, the exception in favor of observation sentences shows, all by itself, that Quine can't unreservedly suppose that the argument from the Q/D thesis and Peirce's thesis to semantic holism is sound. (For some of Quine's recent views on this issue, see his “Reply to Hilary Putnam.”)
12. Among those who reject the Q/D thesis are Dummett, “What is a theory of meaning?”; Glymour, *Theory and Evidence*; Grunbaum, “The falsifiability of theories: total or partial? A contemporary evaluation of

the Duhem–Quine thesis.” Among those who reject verificationism are all the informational semanticists and all the semantic eliminativists (including, from time to time, Quine).

13. It's a vexed question, and one we do not wish to enter into, which (if any) of these options Quine himself intends. (A footnote in his most recent book (*Pursuit of Truth*, p. 78) suggests that it is probably the third.) In any event, the argument we're running is not primarily intended to be *ad hominem*; we claim not just that there is no reading of “statement” that Quine could accept, consistent with his accepting verificationism and Q/D, but that there is no reading of “statement” that makes the putative argument for holism coherent.

We remark in passing that, in the many conversations about these matters that we've had with colleagues, we've been told repeatedly that only one of these options could conceivably have been Quine's intention – indeed, that there's only one worth discussing. Our interlocutors, however, have been divided about equally as to *which* one. In the circumstances, we thought we'd better discuss all three.

14. This is what Grunbaum calls the “trivial” reading of the Q/D thesis (“The falsifiability of theories”).
15. The issue is not, of course, whether it's possible to have a “syntactic” treatment of inductive logic. The question such a treatment would answer is not whether it's *statements* that enter into confirmation relations; it's whether, when they do so, it's in virtue of their *form*.
16. We're grateful to Georges Rey for suggesting another way of reading Quine's use of “statement” which construes statements as formulas but which nevertheless makes a kind of semantic holism come out true. According to this reading, (1) “statement” means *formula*; (2) Peirce's thesis identifies the meaning of a statement (not with its *means* of confirmation but) with its empirical content – namely, with the set of observation sentences that (dis)confirm it (an observation sentence is a formula that is confirmed by/conditioned to proximal stimulations); (3) the Q/D thesis holds as usual. Since meaning is identified with observational consequences and since, according to the Q/D thesis, observational consequences are things that only *total theories* can have, the upshot is that (excepting observation sentences) no formula has a meaning outside a total theory.

This would be a pretty wild sort of semantics. For example, since the meaning of every sentence is the observational consequences of the total theory in which it is embedded, it follows that every sentence in a theory has the same meaning as every other sentence in that theory. Immediate consequences are that no theory can entail a contingent hypothetical, that if a disjunctive statement is true (/false), then both its disjuncts are true

(//false), and so forth. Further, every statement in a theory translates every statement in any empirically equivalent theory (and there are *no* translation relations among statements in theories that aren't empirically equivalent). Perhaps a Quinean might be prepared to accept all this, with the remark "So much the worse for people who insist on having a semantics for sentences." However, there remains what seems to us a crucial difficulty: namely, that the Q/D thesis is trivialized. All that you need to hold onto, in order to hold onto a statement under theory change, is its pronunciation.

17. If you read "concept" the way psychologists usually do, so that concepts are mental particulars, then the first premise of the argument begs the question that's at issue.
18. This is, by the way, the motivation for Jerry Katz's stipulation that if an implication relation among concepts is *semantic* (as opposed to merely *necessary*), then the implied concept must be literally a *constituent* of the implying concept. (See, for example, Katz, *Semantic Theory*.) The effect of this stipulation is to enforce the principle that if one concept analytically implies another, you can't *entertain* (acquire/grasp) the first concept without entertaining (and so on) the second.
In consequence, Katz can hold that it is literally incoherent to ascribe to Jones the belief that some cats purr while denying that Jones believes that some animals do. (We stand neutral as to whether it is well advised of Katz to hold this; whether, for example, it's a view that can be squared with Mates's examples. We also stand neutral on (what appears to be) Katz's assumption that you can't grasp a concept unless you grasp its internal structure.)
19. This bears emphasis because, in later writings, Quine does appear to contemplate a trans-linguistic (trans-theoretic) notion of statement identity (albeit one that is relativized to a choice of "analytic hypotheses"). Whatever other merits this maneuver may have, it doesn't help with the problem of reconciling Peirce's thesis with the Q/D thesis.
20. It doesn't matter, for the discussion that follows, how you construe "conditions of semantic evaluability." For example, they might be satisfaction conditions. Or it might be in Quine's spirit to identify a statement with an ordered pair of a formula and a complex of community-wide speech dispositions of the sort that can be exhibited in publicly observable responses to publicly observable stimuli. In effect, a statement is then a formula together with an inductively certifiable causal regularity in the utterances of the formula. This sort of construal would fit with many of the things that Quine says about translation.
21. More precisely, since every statement must bear to some extent on the level of confirmation of every other (because of simplicity considerations

and the like; see n. 9), what's being said to be an a posteriori issue is *how much* the truth of statement A matters to the truth of statement B.

22. We are grateful to Georges Rey for comments that prompted the two previous paragraphs.
23. This same assumption is implicit in Quine's notorious tendency to identify *languages* with *theories*. One would have thought that, whereas theories are collections of things that have truth values (say, propositions), languages are collections of things that have meanings (say, formulas). But propositions don't *have* meanings because they *are* meanings. So how could theories and languages be the same things? (For some of Quine's views on this issue, see his "Reply to Chomsky.")
24. Since this diagnosis patently begs the question against Quine's skepticism about propositions, we want to emphasize that our case that there is something wrong with the "Two dogmas" argument for semantic holism is internal and does *not* rest on whether we're right about the diagnosis.
25. For example, one might think that the way out of Putnam's Twin worries is to take a "two-factor" view of content, with the "narrow" or "conceptual role" factor being, in effect, verificationist. Quite a lot of people have suggested something of this sort (for example, Block, Field, Loar, Lycan, McGinn, and Putnam). See chapter 6.
26. It's well to notice that even if there is this *ontological* dependence of the syntactic properties of words on the syntactic properties of sentences, no corresponding *epistemological* dependence would follow. In particular, it wouldn't follow that you can't know what syntactic role W has in S unless you know the syntactic structural description of S. It's important that this doesn't follow, because it seems clear that the compositionality of the semantics of a language typically depends on the fact that the syntactic structural description of a sentence is a function of the syntactic structural descriptions of its lexical constituents. It presumably follows that the syntactic analysis of the lexical items must be recoverable ("up to ambiguity") without the prior recovery of the syntax of the sentence. All the grammars and parsing theories we've come across take it for granted that this is so.
The compositionality of syntax is, of course, compatible with there being syntactically ambiguous lexical items. Syntactic compositionality requires only that the entire *range* of possible syntactic analyses of a lexical item be epistemically accessible in each context in which the item occurs. This is what linguists mean when they say that the lexicon of a language specifies the syntactic analysis of words "up to ambiguity."
27. It's a rather nice question whether Russell's treatment of "the F that is G" defines "the" with respect to the semantic or the syntactic properties of its context. What one says about this depends a lot on what one thinks

logical syntax is, a topic that we do not propose to broach here. Suffice it to say that the view that logical syntax really is *syntax* strikes us as perfectly respectable.

28. So, for example, we won't discuss "molecularist" theories of meaning (like Dummett's) which rely on the *a/s* distinction in order to steer between atomism and holism. See chapter 1.
29. Nihilism, and *not* semantic holism, is the view that most people take Quine to hold; and perhaps he does. He says in *Theories and Things*: "Meaning, like thought and belief, is a worthy object of philosophical and scientific clarification and analysis, and like them it is ill suited as an instrument of philosophical and scientific clarification and analysis" (p. 185). Our point is just that he wouldn't be justified in drawing nihilistic conclusions on the basis of the sorts of arguments that are on offer in "Two dogmas" *even if* those arguments are sound against the *a/s* distinction as Quine construes it.
30. He did argue, plausibly in our view, that you can't reconstruct analyticity by appeal to a *prioricity*; that is, that if there *are* analytic truths, we don't know them *a priori*. But this epistemological prohibition needn't imply a prohibition against semantics. Notice that although the kind of Skinnerian theory just sketched can ground a notion of analyticity, it isn't at all committed to the claim that there can be *a priori* knowledge of the analyticities it grounds.

We are grateful to Paul Boghossian for emphasizing the difference between denying that there is an *a/s* distinction and denying that there is an *a/s* distinction grounded in an epistemological property (e.g., *aprioricity*); and also for emphasizing that *any* meaning Realist is sure to endorse *some* notion of content identity. Unlike us, however, Boghossian denies that the arguments in "Two dogmas" undermine an epistemological *a/s* distinction.

NOTES TO CHAPTER 3

1. It's easy enough to run the arguments in the opposite direction and deduce charity from holism. That is, to assume the kind of semantics in which the meaning of a predicate is constituted anatomically, by its role in an inferential network, and then deduce that some of the inferences in which the predicate is implicated – namely, the semantically constitutive ones – must be sound. This is the tactic pursued by Vermazen. Given the assumption that "a predicate gets its meaning by its place in a network of predications corresponding to the user's network of beliefs" ("General beliefs and the principles of charity," p. 114), it's going to follow that "no matter whether we translate the speaker's sentences about individuals charitably, each time we use one of our predicates in translating such a

sentence, we are attributing to him one part of the central network of our own system of beliefs" (*ibid.*, p. 116). Which is to say that we are so translating that many of the informant's general beliefs come out true by our standards.

This argument that network semantics entails charity towards general beliefs is, of course, of no interest for our present purposes, since it starts out by begging the question against the view that the semantic properties of predicates are punctate.

2. It's not actually clear that this second consideration is very forceful. True, a radical interpreter (unlike a bilingual one) is *ipso facto* unable to intuit translation relations between the native language and his home language. But it doesn't follow that it "pre-empts the point of radical interpretation" to assume that he can have *evidence* that certain of the native's sentences translate certain sentences of his own; that would depend on whether the *evidence* for translation must inevitably be couched in semantic terms, not on whether the *success criterion* for translation is couched in these terms. So, pending further argument, there's no obvious reason why accepting Tarski's view of success (adopting Convention T) would beg the question about radical interpretation. Similar remarks would hold, *mutatis mutandis*, if success were defined by an adequacy condition which requires the right-hand side of a T-sentence to be synonymous with, to mean the same as, and so forth, the quoted object language formula on the left.

We're inclined to think that if there is a serious reason for not equating success with material adequacy, it must be the Quinean objection that unexplicated semantic notions like *translation* are not at the disposal of a philosophical account of meaning. Let's suppose this is accepted for the sake of argument. Then "successful truth theory" cannot be elucidated in terms of "materially adequate truth theory."

3. How many of the intuitive semantic relations *is* Davidson requiring a T-theory to "reconstruct"? Equivalently, to what extent is his semantics revisionist? It seems clear that the minimum he wants of the correct truth theory is that it should display whatever information is required to "interpret" the speaker's utterances. That will do for our purposes.
4. The issue *isn't* underdetermination. Even if there are many equally correct but nonequivalent meaning theories, we surely can't allow that a semantics for English should assign all the true sentences to $2+2=4$ and all the false ones to $2+2=5$. But that would be consonant with construing success as extensional adequacy.
5. It's always open to someone to say that, contrary to appearances, child 2's language "really does" have compositional structure, since, by assumption, child 2 can say both that this is snow and that snow is white. But this move is question begging in the present context. If compositional-

HOLISM

A shopper's guide

*Jerry Fodor and
Ernest Lepore*



BLACKWELL
Oxford UK & Cambridge USA