

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's 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, nebulae, 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 è 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 H_2O ," "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, believing *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

INTRODUCTION

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*.

A GEOGRAPHY OF THE ISSUES

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."

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 atomism, then weak atomism 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* atomism; that is, it's not one to which a *weak* atomist is entitled.

The sum and substance of this is that strong atomism 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 atomism 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 atomism 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 atomism 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 –

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.

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 property, 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

NOTES TO CHAPTER 1 (pp. 26–32)

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.

NOTES TO CHAPTER 2 (pp. 33–39)

(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,