5 Holism and Animal Minds

David Finkelstein

1

It's not unusual for influential philosophers to make statements about the minds of nonlinguistic creatures that to a nonphilosopher would seem not merely implausible, but plainly and obviously false. For example, in his *Philosophy and the Mirror of Nature*, Richard Rorty (1979, 187) defends the claim "that knowledge, awareness, concepts, . . . all descend on the shoulders of the bright child somewhere around the age of four, without having existed in even the most primitive form hitherto." According to Rorty, a creature without language—whether human or nonhuman—is not aware of *anything*, and, contrary to what you might imagine when he climbs into bed with you each night, your dog has never once even noticed you.

In arguing for this thesis, Rorty takes himself to be following the lead of Wilfrid Sellars. He quotes approvingly a passage from "Empiricism and the Philosophy of Mind" in which Sellars (1997, 63) describes his own position as one "according to which *all* awareness of *sorts, resemblances, facts*, etc., in short all awareness of abstract entities—indeed, all awareness even of particulars—is a linguistic affair." According to Rorty, Sellars finds an innocent way to allow that brutes may be said to enjoy a *kind* of awareness. As Rorty (1979, 182) reads him, Sellars distinguishes "between awareness-as-discriminative behavior and awareness as . . . being 'in the logical space of reasons, of justifying what one says.'" He goes on: "Awareness in the first sense is manifested by rats and amoebas and computers; it is simply reliable signaling. Awareness in the second sense is manifested only by beings whose behavior we construe as the utterance of sentences with the intention of justifying the utterance of other sentences" (ibid., 182). According to Rorty's Sellars—and Rorty himself—we can, if we like, speak of, for example, your dog's becoming "aware" of you. But this kind of "awareness" (it positively cries out for scare quotes) comes to nothing more than the "awareness" that a thermostat has of the ambient temperature.[1] Thus, while we may

choose to use the word "aware" where all that's at issue is reliable signaling, we should not imagine anything like a continuum of kinds of awareness, with the awareness enjoyed by a dog or a young child lying substantially closer to our sort of awareness than that exhibited by thermostats and waffle irons. We adult humans are conceptually cognizant of objects and states of affairs; dogs and two-year-old children merely respond differentially to stimuli in their environments.

Like Rorty, Donald Davidson holds strikingly skeptical views about the minds of brutes. In his "Rational Animals," he quotes the following from a paper by Norman Malcolm:

Suppose our dog is chasing the neighbor's cat. The latter runs full tilt toward the oak tree, but suddenly swerves at the last moment and disappears up a nearby maple. The dog doesn't see this maneuver and on arriving at the oak tree he rears up on his hind feet, paws at the trunk as if trying to scale it, and barks excitedly into the branches above. We who observe this whole episode from a window say, "He thinks the cat went up that oak tree." (Malcolm 1977, 49; quoted on p. 97 of Davidson 2001a)

Davidson argues that we who say that our dog "thinks the cat went up that oak tree" are mistaken, for a creature without language can have no beliefs; it thinks nothing. His commitment to this claim should be understood as connected to his holism concerning propositional attitudes. Throughout his career, Davidson argues that any belief owes its content to its location in a rich, open-ended network of (among other things) beliefs. Now, why can't your dog share your opinion that a cat has run up an oak tree? Davidson has much to say in reply to this question, some of which we'll examine later, but the following short quotation suggests why it is that, for him, holism about beliefs leads to skepticism about attributing any beliefs to brutes:

[C]an the dog believe of an object that it is a tree? This would seem impossible unless we suppose the dog has many general beliefs about trees: that they are growing things, that they have leaves or needles, that they burn. There is no fixed list of things someone with the concept of a tree must believe, but without many general beliefs, there would be no reason to identify a belief as a belief about a tree, much less an oak tree. Similar considerations apply to the dog's supposed thinking about the cat. (Davidson 2001a, 98)

In order to believe that something is a tree, a creature must know what a tree *is*; it must have the *concept* of a tree. And in order to have the concept of a tree, it must believe quite a number of things about trees, which in turn entails that it have quite a number of other concepts (perhaps *leaf, needle, burning, growth, plant, ground*, and *wood*) whose acquisition requires still *more* beliefs and *more* concepts. In this vein, Sellars argues that the concept of something's *looking green* presupposes not only the

concept of something's *being green*, but also concepts as sophisticated as *standard light-ing conditions*, and he goes on to remark that "there is an important sense in which one has *no* concept pertaining to the observable properties of physical objects in Space and Time unless one has all of them" (Sellars 1997, 44–45).[2]

Thus, a commitment to holism about beliefs and concepts might lead one to con-clude that in order for an ascription that takes the form, "He thinks the cat went up that oak tree," to be true of some creature, the creature in question would need to have a conceptual repertoire about as rich as yours and mine. Given this conclusion, if I were to say of Malcolm's dog, "He thinks the cat went up that oak tree," I would be guilty of an anthropomorphic mistake.

Wittgenstein would appear to be *some* kind of holist about quite a range of things. In his late writings, he seems always to be pointing out how one thing or another—an expectation, a meaningful word, a coronation, an ostensive definition—is what it is only thanks to the "pattern of life" of which it is a part. I have in mind remarks like the following:

Could someone have a feeling of ardent love or hope for the space of one second—*no matter what* preceded or followed this second?—What is happening now has significance—in these surroundings. The surroundings give it its importance. (Wittgenstein 1953, §583)

"Grief" describes a pattern which recurs, with different variations, in the weave of our life. If a man's bodily expression of sorrow and of joy alternated, say with the ticking of a clock, here we should not have the characteristic formation of the pattern of sorrow or of the pattern of joy. (Ibid., p. 174)

When would we say of a child, for instance, that it is pretending? What all must it be able to do for us to say that?

Only when there is a relatively complicated pattern of life do we speak of pretence. (Wittgen-stein 1992, p. 40)

I can't see that Davidson or Rorty would find anything in these passages to complain about. Indeed, it's hard not to view *their* commitments to holism as essentially post-Wittgensteinian. Yet part 2 of the *Investigations* begins as follows:

One can imagine an animal angry, frightened, unhappy, happy, startled. But hopeful? And why not?

A dog believes his master is at the door. But can he also believe his master will come the day after tomorrow? (Wittgenstein 1953, p. 174)

Notice what we *don't* find here. We don't find Wittgenstein saying anything like this: "Since beliefs are intelligible only as situated in the weave of our life, it's a mistake to

think of dogs as having any." Or like this: "While it's not *wrong* to describe a dog as 'believing' that his master is at the door, this way of speaking involves a sort of courtesy that we extend to some selective responders and not others—to dogs and horses but not to thermostats and burglar alarms."[3] Wittgenstein's holism doesn't lead him to the jarring conclusions about animals that we find in Rorty's and Davidson's writings. To understand why this is, we need to come to grips with (what turn out to be) profound differences between Wittgenstein's holism and the sorts of holism that tend to be defended (or attacked) by more contemporary philosophers such as Davidson and Rorty. One of my aims in this essay is to help bring these differences into view.

2

Let's bracket Wittgenstein's holism for the moment; we'll return to it soon enough. Before Wittgenstein was introduced into the discussion, we saw the beginnings of an argument that leads from a commitment to holism about beliefs to what we might regard as a skeptical conclusion—that brutes don't have any beliefs. Such arguments might be, and sometimes are, understood as giving us a reason to *reject* holism about beliefs. Thus, someone might say: "What these arguments really show, or start to show, is *not* that brutes lack beliefs, but that mental holism has absurd consequences. Not only is the holist unable to understand how my dog and I can share a belief (e.g., that some cat is up some tree); the holist can't make sense even of my sharing a belief with *you* unless you share *all* of my beliefs." This complaint about mental holism—that it renders the very notion of agreement between believers problematic—has been articulated by quite a number of philosophers. The most sophisticated version of the argument of which I'm aware is due to Jerry Fodor and Ernest Lepore.

At the start of their book, *Holism: A Shopper's Guide*, Fodor and Lepore (1992, 2) define a holistic property as one such that if one thing has it, then many things must.[4] In their opening chapter, they present an argument that, in one form or another, seems to show that a variety of properties cannot be holistic. Fodor and Lepore set out several versions of the argument. The version most relevant for our purposes goes as follows. They define T* as the property "a belief has iff it expresses a proposition that is the content of some belief of mine" (ibid., 13). Let the word "mine" in this definition index me (not you); then, a token belief of yours may be said have the property T* if and only if it has the same content as some token belief of mine. In other words, if you and I may rightly be said to *share* a belief, then some belief of *yours* has the property T*. Now Fodor and Lepore (1992, 14) ask us to suppose that this property is holistic:

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 [T*]. It would follow that not more than one of us ever has thoughts about color or thoughts about red.

The point could be put as follows: According to, say, the kind of holism that Sellars or Davidson defends, you cannot share only one belief with me; in order to share one, you must share many. But, obviously, you and I don't share *all* our beliefs; we have "widely different belief systems." So, given the truth of this sort of holism, what insures that we share *enough* beliefs for us to share even one?

It is important that we appreciate just what kind of worry is being expressed here. This requires that we not be puzzled by the fact that, in the first sentence of the block quotation above, Fodor and Lepore use the expression "it may well turn out." Their point is not that we should reject holism about belief content because it *entails* that you and I cannot share any beliefs. (It doesn't—or, anyway, it needn't—entail any such thing.)[5] Nor is the point that the truth of this sort of holism would make it too *difficult* for you and me to share a belief:

The problem isn't, notice, that if holism is true, then the conditions for belief identity [i.e., the conditions that would have to be met in order for some belief of yours and some belief of mine to be of the same semantic type—to have the same content] are hard to meet; it's that, if holism is true, then the notion of "tokens of the same type of belief" 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. (Ibid., 19)

A holist about belief content will likely maintain that in order for you to share my belief that, say, cardinals are red, you must share my concept *red*, which in turn entails that you share many beliefs with me having to do with, among other things, colors. But precisely *which* beliefs must you and I share in order for us to share the concept *red*? Fodor and Lepore (1992, 21) ask, "[D]oes believing that Mars is red count more or less for having the concept *red* than believing that tomatoes are?"[6] They complain that holism offers no "principled answer" (21) to such a question. Holism yields no definition of belief-type identity, so no real account of what it takes for two people to share a belief (except in the event that two people share *all* their beliefs). Thus, for all that the holist says, it *could* be—"it may well turn out"—that you and I don't (can't) share a single belief because our networks of belief are insufficiently alike. As I am reading Fodor and Lepore, they think that holism about belief content should be rejected because it ought to rule this out as a possibility, and it fails to.

Notice that this objection to holism presupposes a particular picture of what an account of belief content ought to accomplish.[7] If one wanted to foreground this

picture more than Fodor and Lepore do, one might say something like the following: "Except perhaps when we are doing philosophy, we take it for granted that people often agree about things—that, for example, you and I can, and do, share some beliefs. To say that we take this for granted is not to deny that on some occasion when you and I appear to agree, reasonable doubts may arise about whether we do in fact agree. But we take ourselves to have procedures for determining—in many cases, at least— whether people agree or disagree, and we simply don't question the truism that people do often manage to share beliefs. Now, a satisfactory philosophical account of belief content would, as it were, underwrite or vindicate this truism by specifying—or, anyway, putting us in a position to specify—just what's involved in two people's having the same belief. ('[T]he notion of "tokens of the same type of belief"' should be 'defined' [ibid., 19].) In other words, a satisfactory account of belief content would be, or would yield, an account of *shared* belief. As such, it would enable us to see that the conditions that must be met in order for two people to share a belief are, in fact, often met."

3

We might extract from the discussion thus far two claims about mental holism that merit further consideration. Davidson and Rorty would endorse the first claim. Fodor and Lepore assume the second:

(1)   Holism about beliefs commits one to—or, anyway, suggests compelling reasons for—rejecting the idea that nonlinguistic animals have beliefs.

(2)   A satisfactory philosophical account of belief content would yield an account of shared belief. Thus a satisfactory mental holism would underwrite or vindicate a pair of truisms—that we often agree in our beliefs, and that we are, by and large, correct in our assessments of when we agree.

Either of these claims could be embraced by either a mental holist or an opponent of holism. Someone might accept (1) and count it as a reason for rejecting mental holism. And a holist might grant (2) to Fodor and Lepore, while maintaining (against them) that a satisfactory mental holism has been, or could be, set out.[8]

Here's where I mean to go in the remainder of this essay. In what immediately follows, I'll discuss a lesson that Cora Diamond draws from Wittgenstein. We'll see that if we accept this lesson, then we cannot, with Fodor and Lepore, assume claim (2). Indeed, I'll argue, in §4, that Wittgenstein's continual calls to think about things (e.g., mental states) holistically are meant to help us get over the attraction we feel to

assumptions of this sort. In §5, I'll return to claim (1), that is, to the idea that mental holism gives us reason to be skeptical about animal minds. I'll discus an argument of Davidson's, one of whose conclusions is that brutes have no propositional attitudes, and I'll try to elucidate the differences between Davidson's sort of mental holism and Wittgenstein's. Finally, in §6, I'll say how I think we should answer a question like, "Do dogs really have beliefs and desires?"

Now, let us turn to Cora Diamond and the lives of the saints.

In her paper "Realism and the Realistic Spirit," Diamond points out that according to the norms that governed medieval hagiography, one can write of a saint, for example, that he tried to hide his stigmata, even if he (the hagiographer) adduces no evidence in support of this conclusion—indeed even if the story he tells suggests that the saint in question *wasn't* trying to hide his stigmata. Following Hippolyte Delehaye, Diamond compares hagiography with painting. Just as one may paint a picture of the Baptism of Jesus that shows "besides one man pouring water on the head of another, a dove above their heads, and above that the head of an old man" (Diamond 1991a, 53) without worrying about whether there were, in fact, doves on hand at the original event, just so, "such sentences as 'St. N hid signs of divine favor' may be put into N's *vita* simply because that is how one describes saints" (ibid., 53). Diamond notes that a writer of today who wants to write an accurate biography of a medieval saint might try to address the question of what is true, and not merely conventional, in the old *vitae*. This question makes sense; we distinguish between what it is proper for a medieval hagiographer to write about a saint, on the one hand, and what is true of said saint, on the other. Our acknowledgement of this distinction expresses a commitment to what Diamond (ibid., 55–56) calls "elementary realism":

[W]hat is conventionally put into a saint's *vita* may by no means be true, and we recognize this when we make use, in judging the truth of what is said, of our techniques for weighing and sifting evidence. The fact that these techniques would not have been of interest to the author of the *vita*, who was not attempting to produce an accurate life, nor to his audience, does not mean that our judgments about what is said in the *vita* are in any way out of order or conceptually confused. In dealing with the material in the *vita*, we may perfectly properly adopt a sort of realism: the existence of rules or conventions concerning what may be *said* and indeed *thought* about a certain matter, here saints' lives, leaves open the question what is *true* about those matters. . . . Our elementary realism (as we may call it) has at its heart that contrast between what is said, adhering to the practices, the conventions, governing the writing of saints' lives, and what the facts are.

As I said, my focus in the present section will be a lesson that Diamond takes from Wittgenstein. The lesson could be put as follows: A way in which we often go wrong

when we are doing philosophy is by confusedly taking ourselves to be asking the sort of question that the contemporary biographer of saints asks when he examines old *vitae*. So a philosopher might say, "I know that there are conventions according to which it's perfectly correct for me sometimes to describe myself or someone else as 'following a rule.' What I'd like to determine is whether we really *do* follow rules— whether what is said in these instances is ever really *true*." Diamond (ibid., 55) writes:

I want to suggest that the philosophical realist attempts to take up a position analogous to that of elementary realism—but confusedly. The philosophical realist's conception of *room for* a position analogous to that of elementary realism: *that* is a fantasy.

What's at issue here has as much to do with justification as with truth. Part of Diamond's point—part of what she's getting at when she suggests that the philosophical realist imagines that there is room to take up a position where there is not— is that in doing philosophy, we sometimes describe ourselves as asking whether some assertion is justified, when we have, as it were, so removed ourselves from the familiar terrain on which we know how to give and ask for reasons in support of such assertions that we don't have any real idea what we are asking.

Toward the end of "Realism and Realistic Spirit," Diamond (ibid., 66) quotes the following from Wittgenstein's *Remarks on the Foundations of Mathematics*:

I can train someone in a uniform activity. E.g. in drawing a line like this with a pencil on paper:

— .. — .. — .. — .. — .. — .. — .. — ..

Now I ask myself, what is it that I want him to do, then? The answer is: He is always to go on as I have shewn him. And what do I really mean by: he is always to go on in that way? The best answer to this that I can give myself, is an example like the one I have just given.
    I would use this example in order to shew him, and also to shew myself, what I mean by uniform.
    We talk and act. That is already presupposed in everything that I am saying. (Wittgenstein 1978, 320–321)

Imagine that I draw a line of dots and dashes and ask someone—a child, let's say—to extend the line, going on in the same way. I watch him, and he seems to have no trouble with the task. Now, in a philosophical mood, I wonder: "Has he really done what I asked him to—gone on 'in the same way'? What justifies me in judging that he has?" Perhaps at first, I'm not sure what to do with these questions; it's not clear to me what their answers should look like. At this point, I might seek a philosophical account of what going on in the same way consists in; or of what the words, "going on in the same way," really mean (or, anyway, of what they mean *here*). Such an account would, I suppose, provide me with a kind of test with which I could assess as

true, or as justified, my judgment that the child has gone on in the same way. So I ask myself: "What did I really mean when I told him to go on in the same way? What, exactly, was I asking him to do?"

Wittgenstein suggests that the best answer I can give here is "an example like the one I have just given," that is, an answer that I might enact by looking again at the pattern of marks I drew and saying to myself, "I meant for him to go on doing *that*." Of course, this isn't the *kind* of answer that I take myself to want. Commenting on this passage from *Remarks on the Foundations of Mathematics*, Diamond (1991a, 68) writes, "I want to explain it—and I do not want, do not think I want, something that would in fact, *does* in fact, *do* to explain to someone how to go on." I seek an explanation of what I mean by "go on in the same way." I'm not interested, however, in how I instruct another human being—in what I actually say and do—when I want him to go on in that way. I imagine myself to be asking a question whose answer cannot lie in such mundane instructions and examples. The sort of account I want, or think I want,[9] is one that would reach *past* what I actually do with, or say to, another person. After all, my instructions and examples succeed, if they do, only because I share sensibilities with the person whom I'm instructing. But a philosophical account of going on in the same way cannot take shared sensibilities for granted; it should provide a standard to which our sensibilities—our intuitive assessments of things—can *measure up*. Diamond (ibid., 68–69) writes: "[T]he idea of a philosophical account of what I really mean by 'he is always to go on in that way' is of an account addressed to someone on whose uptake, on whose response, we are not at all depending." I seek an account, an explanation, of that in virtue of which I'm getting it right when I tell the child that *he's* got it right, that is, when I tell him that he has succeeded in extending the pattern that I began. And I won't be satisfied with an account that sounds like this: "I'm right when I tell him that he's got it right by virtue of my recognizing that he has, in fact, continued the pattern I started." For this "recognizing" must be explained too. How do I do it, and by virtue of what is it genuine recognition, rather than misrecognition, or something altogether noncognitive?

As Diamond reads the passage from *Remarks on the Foundations of Mathematics*, it is aimed at getting us to see that this seeming desire for an explanation or an account is not innocent; we can't just take it for granted that when we ask for a philosophical account or explanation of going on in the same way (or following a rule, or being afraid), we are, in fact, asking for anything.

The demands we make for philosophical explanations come, seem to come, from a position in which we are as it were looking down onto the relation between ourselves and reality, some kind of fact or real possibility. We think we mean something by our questions about it. Our questions

are formed from notions of ordinary life, but the ways we usually ask and answer questions, our practices, our interests, the forms our reasoning and inquiries take, look from such a position to be the "rags"[10] [i.e., details that don't matter]. Our own linguistic constructions, cut free from the constraints of their ordinary functioning, take us in. . . . (Ibid., 69–70)

How can we tell if we are being "taken in" in this way? There is no quick, surefire test. Whether or not we have convinced ourselves that there's something, a kind of "account" or "explanation" that we want, when, in fact, we have no coherent idea of anything we want—this needs to be examined on a case-by-case basis. Still, there is much in Wittgenstein's writings to suggest that we should be especially on guard, that we are especially liable to lapse into this sort of illusion, when what we take ourselves to want is a philosophical account of something belonging to the family of phenomena that comprises: following a rule, continuing a series, meaning or intending one thing rather than another, sharing an intention or a belief, and so on.

4

Toward the end of §1, I noted that Wittgenstein seems to be some kind of holist, and I indicated that one of my goals in the present essay would be to help make clearer just what kind he is. As a step toward this goal, I want to repeat a sentence from early on in §3. I said there: "Part of Diamond's point . . . is that in doing philosophy, we sometimes describe ourselves as asking whether some assertion is justified, when we have, as it were, so removed ourselves from the familiar terrain on which we know how to give and ask for reasons in support of such assertions that we don't have any real idea what we are asking." I want to suggest that Wittgenstein urges us to think about, for example, mental and semantic phenomena holistically—in terms of the "weave" of life and of human sensibility—primarily in order to bring us *back* to the familiar terrain on which we know (among other things) how to give and ask for reasons. The aim is *not* to show us where to find raw materials for an account of, say, rule-following that will finally enable us to underwrite our ordinary assessments of when a person may rightly or justifiably be said to be following a rule. It is, rather, to dissolve the illusion that we want anything, when we tell ourselves we want *that*.

Here, it might be helpful to think about two different ways in which holistic-sounding remarks in the *Investigations* might be understood. Consider, for example, the following remarks:

We ask "What does 'I am frightened' really mean, what am I referring to when I say it?" And of course we find no answer, or one that is inadequate.

The question is: "In what sort of context does it occur?" (Wittgenstein 1953, p. 188)

"Then can whatever I do be brought into accord with the rule?"—Let me ask this: what has the expression of a rule—say a sign-post—got to do with my actions? What sort of connexion is there here?—Well, perhaps this one: I have been trained to react to this sign in a particular way, and now I do so react to it.

But that is only to give a causal connexion; to tell how it has come about that we now go by the sign-post; not what this going-by-the-sign really consists in. On the contrary; I have further indicated that a person goes by a sign-post only in so far as there exists a regular use of sign-posts, a custom. (Ibid., §198)

One reading of these passages might be stated as follows: "In the first passage, the one about fear, Wittgenstein is saying that we tend to go wrong philosophically by looking in the wrong place for an account of fear. The only way to arrive at a satisfactory explanation of what fear is—that is, of what 'I am frightened' really means—is to look to the ways in which what we call 'fear' is situated in people's lives, that is, to fear's surroundings. The point of *Investigations* §198 is similar: if we want to know what going by a rule or a signpost really consists in, we must look to the contexts in which people can and do go by rules and signposts, and this means recognizing that any example of someone's going by a rule or a signpost should be understood as part of something larger, a custom. In both passages, Wittgenstein is suggesting that we can arrive at an adequate account of one thing or another only by viewing it as part of a larger whole."

Here, Wittgenstein is read as what we might call a "metaphysical holist"—someone who thinks that questions like "What does going by a signpost really consist in?" or "What does 'I am frightened' really mean?" may be answered by reference to the contexts in which mental and semantic phenomena occur. But Wittgenstein's holism isn't metaphysical. It could be described as "therapeutic." And the felt need for an account—even a *holistic* account—of fear or rule-following that will underwrite (or vindicate or justify) our ordinary practices of assessing when someone has managed to follow a rule is a symptom of the neurosis that it is meant to treat. Thus, in the first quotation above, when Wittgenstein says, "And of course we find no answer, or one that is inadequate," he's not suggesting that we look somewhere else (to context or custom) for a satisfactory explanation of what "I am frightened" really means. The suggestion is, rather, that when we find ourselves asking, "What does 'I am frightened' really mean?" it is likely that we don't understand our own question—likely that we haven't managed to mean anything by it. Wittgenstein's appeals to context, custom, "the weave of our life," and so forth are meant to facilitate an awareness of this point by reminding us of what it looks like for a real person—someone with a life—to be,

for example, afraid, and of what questions we might sensibly ask about (or of) such a person.[11]

5

Let's review. In §1, I sketched an argument that leads from holism about beliefs to skepticism about a nonlinguistic animal's having any beliefs. In §2, we saw an objection to holism about beliefs: that it renders mysterious not only a *dog's* agreeing with me that some cat is in some tree, but *your* agreeing with me as well. I pointed out that Fodor and Lepore's version of this objection assumes that a holist about beliefs owes us an account of mental content that would underwrite our ordinary judgments concerning both *when* people agree and (even) *whether* people ever do. One moral of §§3 and 4 is that a holist need not accept this assumption about what he owes—needn't accept that the assumption even makes sense. To say that we can see a creature as believing something (anything) only if we view it as having *many* beliefs (and perhaps much else besides) is not to admit that we have any grip on what it would mean to provide an account of mental content that would yield a standard to which all our ordinary judgments and sensibilities having to do with agreement would either measure up, and thereby be shored up, or fail to measure up, and thereby be shown incorrect or confused.

In what's to come, I'll return to my opening theme—to the thought that holism about beliefs and concepts tends toward startlingly skeptical conclusions concerning the minds of nonlinguistic animals. You may recall that I drew on passages from Rorty and Davidson when I first sounded this theme. In the present section, I'm going to focus on an argument of Davidson's that purports to show (among other things) that brutes have no propositional attitudes.[12]

Davidson's argument begins with the identification of a problem:

Some creature is taught, or anyway learns, to respond in a specific way to a stimulus or a class of stimuli. The dog hears a bell and is fed; presently it salivates when it hears the bell. The child babbles, and when it produces a sound like "table" in the evident presence of a table, it is rewarded; the process is repeated and presently the child says "table" in the presence of tables. (Davidson 2001b, 117)

We're inclined to say that the dog is responding to the ringing of the bell and that the child is responding to tables. But are we really justified in so saying?

[T]here is a problem about the stimulus. In the case of the dog, why say the stimulus is the ringing of the bell? Why couldn't it be the vibration of the air close to the ears of the dog—or

even the stimulation of its nerve endings. . . . In fact, if we must choose, it seems that the proximal cause of the behavior has the best claim to be called the stimulus, since the more distant an event is causally from its perceiver, the more chance there is that the causal chain will be broken. Why not say the same about the child: that its responses are not to tables but to patterns of stimulation at its surfaces, since those patterns of stimulation always produce the response, while tables produce it only under favorable conditions? (Ibid., 118)

Although the problem is here introduced in connection with nonlinguistic (or, at least, not fully linguistic) creatures, we might ask something similar about ourselves. Imagine that, on my entering your living room, I say, "I like the new coffee table," whereupon you look over at the table appraisingly and reply, "I loved it when I bought it two weeks ago, but now I can't remember what I thought was so great about it." We might ask: What fixes it that I am (or that you are) responding to a particular distal object rather than, say, the more proximal impact of light on my (or, in your case, your) photoreceptors? After all, my reaction, like that of the dog that Davidson describes, is a link in a long causal chain. Why say that I'm responding to the table rather than, for example, to the stimulation of my photoreceptors?

Davidson aims to make intelligible how it is that a behavioral response can be a response to something determinate, and he means to do this by stating a necessary condition on the possibility of such responses.[13] As a step toward bringing this necessary condition into view, notice that in the living room scenario just described, there are *two* responders. So there is, on the one hand, a causal chain running through the coffee table and the stimulation of my photoreceptors to my behavioral response and, on the other hand, a different causal chain running through the table and the stimulation of your photoreceptors to your behavioral response. The common link in the two chains is the table. Davidson (2001c, 213) writes, "We may think of it as a form of triangulation: each of two people is reacting differentially to sensory stimuli streaming in from a certain direction. Projecting the incoming lines outward, the common cause is at their intersection."

Imagine a line of "stimuli streaming" from the direction of the coffee table through the air in the room to my eyes. Now add a second line of stimuli, one that runs through the table via a different route to your eyes. The two lines meet at the table. Finally, add a third line—running between you and me—that represents our awareness of each other as sharing a stimulus in this way, as responding to the same thing. You, the table, and I form a triangle. According to Davidson (ibid., 212–213), if not for the fact that we lead social lives that position us in many such triangles, we could not respond to, talk about, or think about any determinate thing:

[U]ntil the triangle is completed connecting two creatures and each creature with common features of the world, there can be no answer to the question whether a creature, in discriminating between stimuli, is discriminating between stimuli at the sensory surfaces or somewhere further out, or further in. Without this sharing of reactions to common stimuli, thought and speech would have no particular content—that is no content at all.[14]

Now what about nonlinguistic creatures? Are they capable of triangulating, and so responding to, determinate objects? Prima facie, it doesn't seem as if anything I've said thus far rules out the possibility of, say, two dogs triangulating a chew toy. And Davidson (1999, 41) speaks, in connection with wild chimps, of the possibility of a "primitive triangle of two creatures reacting to a common stimulus and to each other's reactions to that stimulus"—indeed of "repeated reactions to shared situations found similar by two or more creatures each simultaneously observing the other's reactions, these reactions in turn being found similar by each creature to reactions previously observed in the other in the shared situations found similar" (ibid., 41). He goes on, in the next paragraph, as follows:

With a single creature, it is hard to decide what it is reacting to when a stimulus hits. When the frog sticks out its tongue, is the stimulus a fly or the firing of a certain pattern of receptors in the eye? We incline to say the latter. But even if the frog were wiser than it is, and learned to save its fire when the target is a large and distant bird or an airplane, we would be in no position to choose one source of the cause of the reaction over another; the frog would just be habituated to a different, though more complex, pattern of proximal stimuli. The slightly complex social situation I have been postulating eases this problem. The stimulus that matters is the nearest mutual cause of the joint reaction. When the triangle is working normally, the mutual reactions of the two (or, of course, more) creatures *triangulate* the relevant stimulus, locating it in a public space. . . . (Ibid., 41)

Here, Davidson seems to be marking a significant distinction between a "single creature," for example, a frog (even an unnaturally "wise" one) and a social, nonlinguistic creature positioned at one apex of a primitive triangle. When a solitary creature exhibits a bit of behavior—although we might "incline to" pick out this or that stimulus as the one to which the thing is responding—we are really in "no position to choose one source of the cause of the reaction over another." In the case of a social brute, however, "[t]he stimulus that matters is the nearest mutual cause of the joint reaction." Thus, it looks as if Davidson holds that primitive triangulation makes responsiveness to determinate stimuli possible. In the same paper, however, he has this to say about Alex, a famous parrot that seems able to discriminate between presented objects' sharing a color and their sharing a shape:

What is it that tells us that the stimulus (cause) of Alex's "answer" to the question "What's the same?" isn't the activation of certain rods and cones in his eyes, or the firing of certain optic

nerves, or the photons bouncing off surfaces we see as the same color? All of these causes, and endless more, are common to the cases where Alex emitted the sound "Color." We have no grounds for choosing one of these causes over the others. (Ibid., 34)

This passage suggests that without language, there just cannot be any fact of the matter concerning what it is that a creature is responding to.

We might try to reconcile these passages in either of two ways. One option is to say that while Davidson thinks some brutes (e.g., chimps) are capable of primitive triangulation and so of responses to determinate objects, he doesn't think Alex is up to either. A second option is to read Davidson as holding that unless a creature is capable of full, linguistic triangulation, there is no fact of the matter about what it is responding to or, anyway, no fact of the matter that is independent of us. How can this second reading be squared with the first passage above? Focus on the word "decide" as it figures in the opening sentence of that passage. The difficulty about the "single creature" is that "it is hard to decide what it is reacting to"—that is, hard for us to decide—where this is not a matter of our discerning that things are a certain way, independently of us. Later in the passage, when Davidson writes, "The slightly complex social situation I have been postulating eases this problem," he means to say that the presence of a primitive triangle eases our problem of deciding for the creature what it is responding to. There's still no independent fact of the matter about what the thing is responding to, as there is when a person responds to a stimulus.[15]

I won't try to settle this exegetical dilemma. For my purposes, it suffices to point out that Davidson does not believe that a primitive triangle—or a long history with such triangles—gives a creature what it needs in order to *think* about a determinate object, state, or event. So, whatever we say about Davidson on the possibility of a dog's *responding to* a determinate thing, it's clear that he does not think that a dog can have, say, a *belief about* a determinate thing. And (at least, for Davidson) this means that a dog cannot have beliefs at all: "Before triangulation can generate thought, the base line between creatures that observe that they share stimuli must . . . be strengthened to include linguistic communication" (Davidson 1999, 42). In order for two creatures, A and B, to participate jointly in the sort of triangle that makes thought possible, A must recognize that B is responding to the same thing that he, A, is responding to. And B must likewise recognize that she and A are responding to the same object. According to Davidson (2001b, 121), this mutual recognition entails that two creatures be capable of rich communication, indeed of *linguistic* communication:

[T]o have the concept of a table or a bell is to recognize the existence of a triangle, one apex of which is oneself, the second apex another creature similar to oneself, and the third an object (table or bell) located in a space thus made common.

The only way of knowing that the second apex of the triangle—the second creature or person—is reacting to the same object as oneself is to know that the other person has the same object in mind. But then the second person must also know that the first person constitutes an apex of the same triangle another apex of which the second person occupies. For two people to know of each other that they are so related, that their thoughts are so related, requires that they be in communication. Each of them must speak to the other and be understood by the other.

Davidson holds that in order to have a concept, a creature must recognize itself and another as reacting to the same stimulus, and this, he thinks, requires linguistic communication. So, as "[t]here is no distinction between having concepts and having propositional attitudes" (Davidson 2004, 137), a nonlinguistic creature can have no propositional attitudes.

At the end of §1, I said that a goal of the present essay is to shed light on the differences between the kind of holism to which Wittgenstein is committed and the kinds that are liable to figure in more contemporary philosophy of mind. I should, therefore, point out that (what we might call) the triangulation theory represents a commitment to a kind of mental holism. According to Davidson, in order for a creature to have a belief about, say, a stick in its path, it must have, as well, beliefs about other creatures and other minds. Davidson (1980, 221) has long held that "[t]here is no assigning beliefs to a person one by one," and the triangulation theory commits him to this claim once again. The theory is holistic in another sense as well, for it tells us that in order for a creature to respond to, and have attitudes about, determinate stimuli, *it* must be a part of something larger, a community of at least two coresponders: "If we consider a single creature by itself, its responses, no matter how complex, cannot show that it is reacting to, or thinking about, events a certain distance away rather than, say, on its skin" (Davidson 2001b, 119). This sounds, in some respects, like something Wittgenstein might say.[16]

But this isn't Wittgenstein's holism. To help make the differences apparent, let us consider a mundane story that illustrates how, in a nonphilosophical context, questions might be raised and settled about the responses and attitudes of nonlinguistic animals. Imagine that I'm out walking my dog, Kita. She is on lead, at my left side. A mailman emerges from a building to our left and walks briskly in our direction. As he approaches, Kita lunges past the front of my body, jerking me around clockwise and to the right. I try to calm her down, the mailman walks past, and a neighbor remarks, "Your dog doesn't want to be anywhere *near* that mailman." I reply: "Actually, they get along fine. Her reaction wasn't to him. It was to that squirrel." Here, I point to a squirrel in a nearby tree that's looking down at Kita and chittering loudly. "It ran past

us, and she wanted to chase it. Of course, now that she knows there's a squirrel sitting there, I'll have to drag her away from here."

This is how it might look and sound for questions to be settled about a nonlinguistic animal's reactions and attitudes. Davidson, of course, knows this, that is, he knows what ordinary conversations of this sort look and sound like. But he thinks that, viewed from a *philosophical* vantage point, the justification that I offer (or could offer) to my neighbor for the claims that I make about Kita may be seen as taking too much for granted.

What, exactly, do I take for granted that, according to Davidson, ought not to be assumed from a philosophical vantage point? Here's one thing: that Kita, being a dog, is the sort of animal that *hunts*. For hunting involves (or just is) a kind of responsiveness *to* some sorts of things and not others. A question like, "What entitles us to claim that Kita hunts *squirrels* and *rabbits* rather than, say, photons or photoreceptors?" is not, I take it, one that anybody would try to answer seriously. ("Well, as for photons, they're just too *fast* to chase down.") Thus, in order for Davidson's problem— we might call it the problem of "stimulus determination"—to get a grip on us, we have to view animals and human beings merely as potential *responders*, not as *hunters* or as *chasers*. We mustn't help ourselves to the concept of hunting or of chasing. And the point extends beyond *hunting* and *chasing* to, for example, such concepts as *fleeing*, *eating*, *mating*, *feeding*, *caring for*, and *playing*. ("Why should we describe that bear as caring for her cub, rather than for her sensory surfaces?" is not a serious question either.) Davidson insists, in effect, on our setting aside all such concepts when we ask ourselves which stimulus, if any, a creature might be responding to. Here, then, is a way to characterize one thing that must not be taken for granted if the problem of stimulus determination—the problem to which triangulation is meant as a solution— is to seem gripping in the first place: that people and animals have *lives*.[17]

In §1, I noted that Wittgenstein's holism expresses itself in his continually urging us to think about mental and semantic phenomena as they figure in what he calls "the weave of our life." He believes that philosophical confusion about meaning and intentionality (as well as about sensations and privacy) typically results from viewing signs and behavior as cut off from life,[18] whereas Davidson's triangulation theory *requires*, in effect, that we consider the behavior of animals and human beings *as lifeless*.

So, that's one significant difference between Davidson's holism and Wittgenstein's. In what follows, I'll try to spell out another.

Davidson's triangulation theory offers, or seems to offer, a philosophical vantage point from which we can judge—as properly justified or not—our ordinary procedures for

determining whether this or that creature is responding to, or having thoughts about, this or that object. The theory appears to show that a great deal of what we say about animals, even animals with whom we share our daily lives, is not really true or justified. It seems to give us a kind of *test* of our prephilosophical assessments of, and sensibilities having to do with, the ascription of responsiveness and intentionality to the animals and human beings around us. We usually pass the test when we speak about adult human beings. We tend to fail in our dealings with brutes and young children.

Here, I want to remind you of an example that was discussed in §3: I draw a lines of dots and dashes on a piece of paper and ask a child to extend it, going on in the same way. On observing the child continue the pattern, I wonder whether I'm really justified in judging that he's done what I asked him to do. I imagine that a philosophical account of *going on in the same way* might underwrite my judgment, that such an account might serve as a standard against which my assessment of the child's behavior would either measure up or fail to measure up. The moral of §3's discussion of this example was that in our search for philosophical accounts or explanations that will serve either to underwrite or to undermine our everyday sensibilities, we are inclined toward a kind of nonsense. By trying to take nothing for granted—by trying to step *way back* from our ordinary procedures for assessing and justifying claims having to do with, for example, whether someone has managed to continue a series or follow an instruction—we leave ourselves without the resources we need in order even to make sense of phrases like "continuing a series," "following an instruction," or "underwriting a judgment to the effect that a child has managed to go on in the same way."[19] And rather than come to grips with this fact (i.e., with the fact that we've lost our grip on our own words), we're liable instead to think that, from the privileged vantage point we have attained by stepping back from our ordinary procedures of assessment and justification, we now see that there is really nothing to the idea of continuing a series or following a rule. We're liable, in this way, to confuse a self-inflicted inability to make sense of some semantic or mental phenomenon with grounds for being skeptical about it.

We've seen that Davidson assumes, in effect, that the truth about an animal's state of mind, or lack thereof, comes into view only when we set aside most of what we know about the animal, *qua* animal, and try to view its movements as though they were not the doings of a living creature. One way to criticize Davidson's position is to point out that he's not entitled to this assumption. But if the Wittgensteinian line of thought sketched in §§3 and 4—and recalled in the preceding paragraph—is sound, then perhaps a further criticism can be made here. Perhaps we can say that the "philosophical vantage point" Davidson would have us attain turns out to be only the illusion of a vantage

point. And if *this* is correct, then an important difference between Davidson's holism and Wittgenstein's can be put as follows: the latter is a cure for the former.

6

In §3, I asked you to imagine a philosopher who says, "I know that there are conventions according to which it's perfectly correct for me sometimes to describe myself or someone else as 'following a rule.' What I'd like to determine is whether we really *do* follow rules—whether what is said in these instances is ever really *true*." A lesson that emerged in the ensuing discussion was that we ought not simply to assume that words like these are *always* significant. We should not take it for granted that this speaker has yet found anything to mean by the word "really" in its occurrences here. Now imagine that someone puts the following question to me: "I understand that, by ordinary standards, you're entitled to say that Kita wanted to chase that squirrel and that she knew it was in that tree. What I'd like to determine is this: Did she *really* want to chase it? And did she *really* know it was in that tree?" It seems to me that the best answer one could give to such an interlocutor would not be "Yes" or "No," but, as a start, something more like this: "I don't think I understand your question."[20]

Appendix

What about the argument for skepticism about animal minds that was (barely) touched on in §1? Perhaps it seemed more modest than the triangulation argument discussed in §5 and, for this reason, more convincing. In what follows, I'll walk you through Davidson's version of the §1 argument, and I'll say where I think it falls short.

Davidson (2001a, 97) says, "One way of telling that we are attributing a propositional attitude is by noting that the sentences we use to do the attributing may change from true to false if, in the words that pick out the object of the attitudes, we substitute for some referring expression another expression that refers to the same thing." Propositional attitude ascriptions exhibit semantic opacity. But now, Davidson claims, opacity is absent (or, anyway, not clearly enough present) when we try to ascribe beliefs to brutes. Thus, concerning the dog described by Malcolm, Davidson (ibid., 97) writes:

[H]ow about the dog's supposed belief that the cat went up that oak tree? That oak tree, as it happens, is the oldest tree in sight. Does the dog think that the cat went up the oldest tree in sight? Or that the cat went up the same tree it went up the last time the dog chased it? It is hard to make sense of these questions.

The argument here might be schematized as follows:

(1) A distinguishing feature of propositional attitude attributions is that they exhibit semantic opacity: the truth of such an attribution may be affected by how some object is referred to.

(2) When we attribute attitudes—beliefs, say—to brutes, the attributions do not exhibit opacity. So we don't find that a belief attribution that picks out a tree *in one way* (as, e.g., "Mr. Malcolm's oak tree") is true, while an attribution that differs only in that it picks out the same tree in a different way (as "the oldest tree in sight") is false.

(3) Brutes don't really have propositional attitudes.

Let's call this the opacity argument.

As I've schematized the argument, the gap between steps (2) and (3) seems rather wide. We might try to narrow it by noting the following: Belief attributions exhibit opacity *because* we think of objects in only some of the ways in which they may be thought of. (This, in turn, reflects our finitude. Belief attributions made to God shouldn't exhibit opacity.) If, when we try to ascribe beliefs to dogs, the ascriptions do not exhibit opacity, this suggests (either that those bumper stickers that read "Dog is my co-pilot" are *really* onto something or) that dogs aren't thinking of the objects to which we're referring *in any way at all*. Now, if a dog isn't thinking of, say, a particular tree in any way at all, then (the thought goes) it can't really be believing anything about the tree; it can't really believe, for example, that some cat went up it.

Perhaps the most obvious objection that might be raised against the opacity argument is this: it *does* seem to matter how we refer to objects when we attribute beliefs to animals; such attributions do seem to exhibit (at least some degree of) opacity. Just compare "The dog thinks the cat went up that tree" with "The dog thinks the cat went up the oldest oak tree in Cook County." We can't take the latter seriously (except as a *de re* ascription). This isn't true of the former. We're inclined to think that while a dog cannot think of some object as the oldest tree in Cook County, it can think of something as a tree. Thus, we're inclined to suppose that a belief attribution to Malcolm's dog that referred to some tree as "that tree" might capture, as it were, the dog's *take* on things better than some other attribution that was different only in that it referred to the same tree in another way. If this supposition is correct, then the opacity argument is undermined.

We can think of the following passage, part of which we saw in §1, as representing Davidson's reply to this objection:

In a popular if misleading idiom, the dog must believe, under some description of the tree, that the cat went up the tree. But what kind of description would suit the dog? For example, can the dog believe of an object that it is a tree? This would seem impossible unless we suppose the dog has many general beliefs about trees: that they are growing things, that they have leaves or needles, that they burn. There is no fixed list of things someone with the concept of a tree must believe, but without many general beliefs, there would be no reason to identify a belief as a belief about a tree, much less an oak tree. Similar considerations apply to the dog's supposed thinking about the cat. (Davidson 2001a, 98)

According to Davidson, a dog can no more think of something as a tree than as the oldest tree in sight. Why not? Well, in order for a creature to think of something as a tree, it must have some idea what a tree *is*, and so have "many general beliefs about trees."

We should note that Davidson does not assume that a dog couldn't be supposed to have *any* general beliefs about trees. This would be to beg the question. There are, after all, *some* general facts that we might—until we were *convinced* by Davidson's argument—suppose dogs to know about trees: that squirrels climb them, that cats climb them, that squirrels jump from one to another, that there is often dog urine on them, that leashes can get wrapped around them, that they are outside, and so forth. What Davidson claims is that there are, *as well*, many general facts about trees (e.g., "that they are growing things") that we would expect a normal, adult human being to know, but that a dog could not plausibly be supposed to know or to believe. His next move, which is made only implicitly in the quoted passage, is to claim that although, as just noted, we might suppose, prima facie, that dogs have *some* general beliefs or knowledge about trees, even *given* this supposition, they don't have *enough* general beliefs or knowledge to think of an object as a tree.

The trouble with the opacity argument lies in this last move. For how are we to say how many general beliefs about trees a creature must have in order to think of something as a tree? If we rely merely on our intuitions here, we'll probably say that however many general beliefs a dog has is enough. (After all, the skeptical conclusion isn't intuitively appealing.)

At this point, the opacity argument might seem to stand in need of supplementation, that is, it might seem that Davidson does, after all, need to specify precisely *how much* a creature must believe or know concerning trees in order to be capable of thinking of anything as a tree. An alternative argumentative strategy would be for him to claim that some *particular* belief or piece of knowledge is both necessary for thinking of something as a tree *and* impossible for brutes to acquire. In §5, we saw Davidson arguing for just such a claim. According to (what I there called) the triangulation

theory, a creature cannot think of an object as a tree (or as anything else) unless it knows something that, according to Davidson, no brute can know, namely, that it is one apex of a triangle (whose second apex is a creature similar to itself and whose third is an object "located in a space thus made common" [Davidson 2001b, 121]). Thus, a consideration of the skeptical argument touched on briefly in §1 leads us back to the argument from triangulation discussed at greater length in §5.

Notes

1. I should emphasize that, in this paragraph, I mean to be describing a position held by Rorty and by *Rorty's* Sellars. I'm bracketing the question of whether the skeptical views concerning animals that Rorty takes himself to find in "Empiricism and the Philosophy of Mind" are actually there. It *is* tempting to read EPM-on-animals more or less as Rorty does. But Sellars's subsequent work suggests a different reading (see Sellars 1975, 303–304; and Sellars 1981).

2. Jerry Fodor and Ernest Lepore (whom I'll be discussing in the next section) complain that this sort of holism makes the acquisition of language look impossible:

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? (Fodor and Lepore 1992, 9)

Sellars considers (what amounts to) this question in "Empiricism and the Philosophy of Mind." His answer could be put as follows: A child doesn't learn the meanings of words one by one. Instead he masters "*piecemeal* habits of response" (Sellars 1997, 45) until his behavior comes to look enough like that of his elders that it makes sense to credit him with knowledge of a great *many* words (and facts). Learning a language requires—not a spasm of cognition, but—several years during which a child's behavior comes gradually to resemble that of an adult.

3. After noting that "[b]abies and the more attractive sorts of animal are credited with 'having feelings' rather than . . . 'merely responding to stimuli,'" Rorty (1979, 189–190) suggests that we understand this "as a courtesy extended potential or imagined fellow-speakers of our language."

4. Understood as an indication of how philosophers use the word "holistic," this is a little misleading. According to Fodor and Lepore's definition, Davidson commits himself to holism about beliefs—or, anyway, to the view that the property of being a belief is holistic—when he says that in order for a creature to have even one belief, it must have many beliefs (when he claims, in other words, that for even one of a creature's states to have the property of being a belief, many of them must). But Davidson *also* thinks this: in order for a creature to have even one belief, it must have many beliefs *and* desires; *moreover*, it must engage in many actions. This view of Davidson's represents a commitment to a different kind of holism about beliefs—one that involves a "whole" that comprises more than just a lot of beliefs. I take it that the meaning of "holism," both in and out of philosophy, involves some idea like this: a single part of something can't be understood, or even be what it is, apart from the whole. Thus, a holist about beliefs thinks that a single belief can't be understood apart from *some* larger whole of which it may be understood

to be a part. The larger whole need not be understood as a network of beliefs, however. It might be understood as a network of beliefs, desires, and actions—or, as in Wittgenstein's late writings, something broader still, something more like a life.

5. A surprising number of philosophers—both holists and critics of holism—write as if holism about beliefs does entail that two people cannot share any beliefs unless they share all their beliefs. Thus, having defined "meaning holism" as the view, "roughly, that each representation in a linguistic or mental system depends semantically on every other representation," Eric Lormand (1996, 51) writes: "If meaning holism about a system *S* is true, then a change in meaning of any representation in *S* requires a change in meaning of all representations in *S*." And in the *Routledge Encyclopedia of Philosophy*, Ned Block (1998, 488) says that if "the content of any state depends on all the others, it would be extremely unlikely that any two believers would ever share a state with the same content." Notice that whether or not such statements are true depends entirely on the way in which each element in a system of states or representations is understood to *depend* on the others. There is, after all, more than one kind of dependence. Thus, Henry Jackman (1999, 363) writes:

Holism only requires that the content of any one of one's beliefs *depend upon* or be a *function of* one's other beliefs, and *this* claim need not commit one to the instability thesis [i.e., the thesis that any change in a network of beliefs changes the content of every belief in the network]. After all, one can claim that *A* is a function of *B* without implying that any change in *B* will produce a change in *A*. Consider, for instance, the claim that one's final letter grade in a class is a function of (depends upon) the results of one's exams, quizzes and homeworks. The truth of this claim certainly doesn't entail that no two people could have the same final grade unless they had precisely the same score on all of their homeworks, exams and quizzes.

The point seems hard to miss, and we might wonder why so many philosophers—both holists and critics of holism—appear to take seriously the worry that mental holism entails that two people cannot really agree about anything unless they agree about everything. Block, who is himself a mental holist, recommends that we give up the very idea of agreement and replace it with a graded notion, viz., similarity of content. I would suggest that what's moving Block and others is, ultimately, a metaphilosophical commitment that they share with Fodor and Lepore. I'll say a bit more about this in n. 8.

6. There is a difficulty with the question thus formulated: many philosophers (Fodor and Lepore, among them) would say that someone's believing *either* that Mars is red *or* that tomatoes are red counts *decisively* in favor of ascribing the concept *red* to him. Thus neither of the beliefs mentioned would count either more or less "for having the concept *red*." Fodor and Lepore might, however, have avoided this difficulty by posing a slightly different question, viz., Which counts more *against* the ascription of the concept *red* to someone, that he does *not* believe that Mars is red, or that he does *not* believe that tomatoes are?

7. It may be helpful, as we think about what the objection presupposes, to consider another version of it (one directed at holism about another property): Fodor and Lepore (1992, 12) define R* as the property an expression has "iff it refers to something or other that currently accepted astronomical theories refer to." Now suppose that R* is holistic; suppose, that is, that if one expression in the language of a particular theory refers to something or other that our current

astronomical theories refer to, then many expressions in the language of that theory must also refer to things that our current astronomical theories refer to:

Then it might turn out that no theory could refer to (for example) stars unless it could 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 astronom*ers*) 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*. (Ibid., 12)

Fodor and Lepore think it undeniable that Greek astronomers did, in fact, refer to stars. Still, this argument isn't a reductio ad absurdum. We should reject holism here not because it entails that the Greeks had no views about stars, but because it leaves this open as a possibility. Hence, the passage begins with the words "Then it might turn out . . ." rather than "Then it turns out . . ."

8. Ned Block would reject claim (2). As I indicated in n. 5, he accepts the idea (even though he admits it is "weird-sounding" [Block 1998, 491]) that you and I cannot rightly be said to agree about anything. By his lights, we preserve enough of common sense if we allow that you and I have many beliefs that are *similar* in content. Still, while Block would reject (2), he would, I think, accept a weaker version of it (and this is the metaphilosophical commitment that I mentioned at the end of n. 5):

(2′) A satisfactory philosophical account of belief content would yield an account of shared belief. Thus a satisfactory mental holism would put us in a position to *evaluate* as true or false both the claim that people often agree in their beliefs and that we are, by and large, correct in our assessments of when we agree.

As should become apparent in what follows, I think Wittgenstein would urge us to question (2′) along with (2).

9. The move here from speaking about what I want to what I *think* I want doesn't, after all, accomplish very much. If I don't manage to mean anything by, e.g., "an account of going on in the same way," then it doesn't make sense for me to describe myself *either* as wanting an "account of going on in the same way" *or* as thinking that I do. In writing *about* philosophical nonsense, one sometimes finds that one's own sentences turn out to be nonsense. If one is doing things right, they are ladders that, eventually, could be thrown away.

10. This is a reference to *Investigations* §52:

If I am inclined to suppose that a mouse has come into being by spontaneous generation out of grey rags and dust, I shall do well to examine those rags very closely to see how a mouse may have hidden in them, how it may have got there and so on.

   But if I am convinced that a mouse cannot come into being from these things, then this investigation will perhaps be superfluous. But first we must learn to understand what it is that opposes such an examination of details in philosophy.

11. A reader of *Philosophical Investigations* who does not get the therapeutic aim of Wittgenstein's holism into focus is liable to be struck by some such thought as the following: "While Wittgenstein does point us in the direction of an account of fear (or rule-following or intentionality), he

doesn't seem particularly interested in working out the details and actually providing the account. Indeed, he seems oddly, perhaps even perversely, opposed to *anyone's* providing the account—like a Moses who, having finally led the Jews to a place from which they can *see* the Promised Land, says, 'I hope none of you people are planning to actually go in there.'" Wittgenstein's commentators too often, I think, leave one with this sort of impression of him. Consider, e.g., Robert Brandom. Early on in his *Making It Explicit*, Brandom (1994, 29) claims that (1) if "anything is to be made of the Kantian insight that there is a fundamental normative dimension to the application of concepts," then a (certain sort of) "theory of practices" must be spelled out, and (2) even though Wittgenstein's arguments establish the "criteria of adequacy" (ibid., 30) that such a theory would have to meet, he (Wittgenstein) would urge *against* the pursuit of a theory of practices:

Wittgenstein, the principled theoretical quietist, does not attempt to provide a theory of practices, nor would he endorse the project of doing so. The last thing he thinks we need is more philosophical theories. (Ibid., 29)

Although Brandom here calls Wittgenstein's quietism "principled," he doesn't explain what (presumably misguided) principle underlies it. (Surely, it can't be, merely, that the last thing we need is more theories. Moses: "The last thing we need is more walking.") A bit later in his book, Brandom articulates a twofold explanatory task that he aims to meet in order to entitle himself to the concept of intentional content. The task is, first, "to say what it is to *express* a *propositional* content in general, and then to say what more is required specifically for the content expressed to *represent* something *objective*, in the way that matters for empirical science" (ibid., 75). Brandom goes on:

This is a request that can sensibly be addressed to Wittgenstein, as well. Even his sustained, penetrating discussions do not offer an account of what distinguishes language games within which states and performances acquire specifically *propositional* significances . . . , nor of what distinguishes those within which states and performances acquire specifically *representational* significances. (Ibid.)

Here, Brandom doesn't bother to suggest that Wittgenstein's unwillingness to pursue such accounts is principled.

12. By my count, Davidson sets out three arguments that lead to this conclusion: two in "Rational Animals" (Davidson 2001a; originally published in 1982) and one more—the one that I'm about to describe—in a number of his later papers (several of which I'll cite in what follows). Of the three arguments, he seems most satisfied with this last one. (In the appendix, I'll discuss one of the two earlier arguments.)

13. Obviously, it won't do to state just *any* necessary condition, e.g., that in order for a creature to respond to determinate objects, it must be capable of *movement*. This would be, by Davidson's lights, true, but it wouldn't help to make the possibility of such responses intelligible.

14. Remarks such as this one might suggest that when I'm by myself, I'm unable to think about determinate objects. Of course, this isn't Davidson's view. Rather: my rich history of triangulating with others enables me to think about particular things even when no one else is around. The point can be put in terms of what interpretation requires. If, alone in my own living room, I say to myself, "That table needs to be fixed," I'm best interpreted as responding to, and thinking about, my table—rather than, say, events in my own eyes—thanks to my history of

triangulating with others. If not for this history, however, there wouldn't be sufficient constraint on how best to interpret me for me to be rightly viewed as thinking *about* anything.

15. Bridges (forthcoming) reads Davidson's remarks about Alex as expressing this sort of antirealism about the responses of (even social) brutes. He writes: "[T]here is really nothing about the behavior, considered in and of itself, that licenses viewing it as specially linked to certain of its causes. What explains our favoritism toward these causes is not something that we discover in the behavior, but something we bring to it, namely our 'natural' dispositions to attend to certain kinds of events and ignore others."

16. To make it sound *more* like something Wittgenstein might say, let this imagined creature be supposed to exist for only an instant.

17. In this paragraph, I've been more or less paraphrasing one of the main points in an insightful paper on Davidson by my colleague Jason Bridges. He writes:

To refrain from conceiving an animal as responding to distal events and objects is to refrain from conceiving it as fleeing from anything. Or, by the same token, as playing with, watching over, hiding from, searching for, returning to, threatening, fighting with, attacking, defending territory from, following, stalking, herding, foraging, burying, stealing, greeting, communicating with, grooming, etc. (Bridges forthcoming)

18. Here are three (of many) passages from the *Investigations* that might be considered in connection with this point:

How does it come about this arrow → *points*? Doesn't it seem to carry in it something besides itself?—"No, not the dead line on paper; only the psychical thing, the meaning, can do that."—That is both true and false. The arrow points only in the application that a living being makes of it. (§454)

It is possible to say "I read timidity in this face" but at all events the timidity does not seem to be merely associated, outwardly connected, with the face; but fear is there, alive, in the features. (§537)

Look at a stone and imagine it having sensations.—One says to oneself: How could one so much as get the idea of ascribing a *sensation* to a *thing*? One might as well ascribe it to a number!—And now look at a wriggling fly and at once these difficulties vanish and pain seems able to get a foothold here, where before everything was, so to speak, too smooth for it. (§284)

19. This moral should not be overstated. A lesson I would not want to draw from Wittgenstein (or anyone) is that whenever someone who calls herself a philosopher offers a theory or an account of something suggesting that our everyday sensibilities are, to a certain extent, mistaken—i.e., that our ordinary procedures of assessment and justification ought to be revised— she lapses into nonsense or falls into error. I am confident that Diamond agrees with me about this, and *because* I am, I take a passage like the following to be somewhat misleading:

The sense in which philosophy leaves everything as it is this: philosophy does not put us in a position to justify or criticize what we do by showing that it meets or fails to meet requirements that we lay down in our philosophizing. (Diamond 1991b, 22)

We could, I suppose, stipulate that the word "philosophy" be used such that when we *are* "in a position to justify or criticize what we do by showing that it meets or fails to meet requirements," we don't count as doing philosophy. But I can't see that this is a good idea.

20. I'm grateful to Jason Bridges, James Conant, Cora Diamond, Jay Elliott, Aidan Gray, Erica Holberg, and Thomas Lockhart for helpful comments on drafts of this essay.

Works Cited

Block, N. 1998. "Holism: Mental and Semantic." In E. Craig, ed., *Routledge Encyclopedia of Philosophy*, 488–493. London: Routledge.

Brandom, R. 1994. *Making It Explicit*. Cambridge, Mass.: Harvard University Press.

Bridges, J. "Davidson's Transcendental Externalism." Forthcoming in *Philosophy and Phenomenological Research*.

Davidson, D. 1980. "Mental Events." In *Essays on Actions and Events*. Oxford: Oxford University Press.

Davidson, D. 1999. "Interpretation: Hard in Theory, Easy in Practice." In M. De Caro, ed., *Interpretations and Causes: New Perspectives on Donald Davidson's Philosophy*, 31–44. Dordrecht: Kluwer Academic Publishers.

Davidson, D. 2001a. "Rational Animals." In *Subjective, Intersubjective, Objective*, 95–105. Oxford: Oxford University Press.

Davidson, D. 2001b. "The Second Person." In *Subjective, Intersubjective, Objective*, 107–121. Oxford: Oxford University Press.

Davidson, D. 2001c. "Three Varieties of Knowledge." In *Subjective, Intersubjective, Objective*, 205–220. Oxford: Oxford University Press.

Davidson, D. 2004. "What Thought Requires." In *Problems of Rationality*, 135–149. Oxford: Oxford University Press.

Diamond, C. 1991a. "Realism and the Realistic Spirit." In *The Realistic Spirit*. Cambridge, Mass.: MIT Press.

Diamond, C. 1991b. "Introduction II: Wittgenstein and Metaphysics." In *The Realistic Spirit*, 39–72. Cambridge, Mass.: MIT Press.

Fodor, J., and E. Lepore. 1992. *Holism: A Shopper's Guide*. Oxford: Basil Blackwell.

Jackman, H. 1999. "Moderate Holism." *American Philosophical Quarterly* 36:361–369.

Lormand, E. 1996. "How to Be a Meaning Holist." *Journal of Philosophy* 93:51–73.

Malcolm, N. 1977. "Thoughtless Brutes." In *Thought and Knowledge*, 40–57. Ithaca: Cornell University Press.

McDowell, J. 1998. "Wittgenstein on Following a Rule." In *Mind, Value, and Reality*. Cambridge: Harvard University Press.

Rorty, R. 1979. *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.

Sellars, W. 1975. "The Structure of Knowledge." In H. Castañada, ed., *Action, Knowledge, and Reality: Critical Studies in Honor of Wilfrid Sellars*, 295–347. Indianapolis: Bobbs-Merrill.

Sellars, W. 1981. "Mental Events." *Philosophical Studies* 39:325–345.

Sellars, W. 1997. *Empiricism and the Philosophy of Mind*. Cambridge, Mass.: Harvard University Press.

Wittgenstein, L. 1953. *Philosophical Investigations*. Trans. G. E. M. Anscombe. Oxford: Basil Blackwell.

Wittgenstein, L. 1978. *Remarks on the Foundations of Mathematics*. Ed. G. H. von Wright, R. Rhees, and G. E. M. Anscombe. Oxford: Basil Blackwell.

Wittgenstein, L. 1992. *Last Writings on the Philosophy of Psychology*, vol. 2. Trans. C. G. Luckhardt and M. A. E. Aue. Ed. G. H. von Wright and H. Nyman. Oxford: Blackwell.