



Truth Rules, Hoverflies, and the Kripke-Wittgenstein Paradox

Ruth Garrett Millikan

The Philosophical Review, Vol. 99, No. 3 (Jul., 1990), 323-353.

Stable URL:

<http://links.jstor.org/sici?sici=0031-8108%28199007%2999%3A3%3C323%3ATR%3E2.0.CO%3B2-5>

The Philosophical Review is currently published by Cornell University.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/sageschool.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Truth Rules, Hoverflies, and the Kripke-Wittgenstein Paradox

Ruth Garrett Millikan

[T]he sceptical argument that Kripke attributes to Wittgenstein, and even the 'sceptical solution', are of considerable importance regardless of whether they are clearly Wittgenstein's. The naturalistically inclined philosopher, who rejects Brentano's irreducibility and yet holds intentionality to be an objective feature of our thoughts, owes a solution to the Kripke-Wittgenstein paradox.¹

The challenge is a welcome one. Although I will argue that the Kripke-Wittgenstein paradox is not a problem for naturalists only, I will propose a naturalist solution to it. (Should the Kripke-Wittgenstein paradox prove to be soluble from a naturalist standpoint but intractable from other standpoints, that would, I suppose, constitute an argument for naturalism.) Then I will show that the paradox and its solution have an important consequence for the theories of meaning and truth. The Kripke-Wittgenstein arguments which pose the paradox also put in question Dummett's and Putnam's view of language understanding. From this view it follows that truth rules must be "verificationist rules" that assign assertability conditions to sentences, rather than "realist rules" that assign correspondence truth conditions. The proposed solution to the paradox suggests another view of language understanding, according to which a speaker can express, through his language practice, a grasp of correspondence truth rules. This will block one route of Putnam's famous retreat from realism:

The point is that Dummett and I *agree* that you can't treat understanding a sentence (in general) as knowing its truth conditions; because it then becomes unintelligible what *that* knowledge *in turn* consists in. We both *agree* that the theory of understanding has to be done in a verificationist way . . . conceding that *some* sort of verificationist semantics must be given as our account of understanding. . . . I have

¹Brian Loar, Critical Review of Saul Kripke's *Wittgenstein on Rules and Private Language*, *Noûs* 19 (1985), p. 280.

given Dummett all he needs to demolish metaphysical realism . . . a picture I was wedded to!²

(By “metaphysical realism” Putnam means, roughly, the traditional correspondence theory of truth.) Elsewhere I have argued that the distinction Putnam draws between “metaphysical realism” and “internal realism” is illusory, that naturalist arguments for correspondence truth are, inevitably, arguments for truth as correspondence to theory-independent objects, and that there is nothing incoherent in this notion of correspondence.³ So in giving a naturalist argument to show that grasping correspondence truth rules is no more problematic than grasping verificationist ones, I take myself to be defending the strongest possible kind of correspondence theory of truth and the most flatfooted interpretation possible of the truth-conditions approach to semantics.

I. THE KRIPKE-WITTGENSTEIN PARADOX

The Kripke-Wittgenstein paradox, as Kripke explains it, is an apparent dead end we encounter when trying to explain what it is that constitutes a person’s meaning something by a word. Kripke takes addition as his central example: what constitutes my meaning addition by “plus” or “+”? “Although I myself have computed only finitely many sums in the past, the rule for addition determines my answer for indefinitely many new sums that I have never previously considered. This is the whole point of the notion that in learning to add I grasp a rule: my past intentions regarding addition determine a unique answer for indefinitely many cases in the

²Hilary Putnam, “Realism and Reason,” in *Meaning and the Moral Sciences* (London, England: Routledge and Kegan Paul, 1978). Michael Dummett’s statement is in “What is a Theory of Meaning?” in Samuel Guttenplan, ed., *Mind and Language* (Oxford, England: Clarendon Press of Oxford University Press, 1975), pp. 97–139, and in “What is a Theory of Meaning? (II),” in Gareth Evans and John McDowell, eds., *Truth and Meaning: Essays in Semantics* (Oxford, England: Clarendon Press of Oxford University Press, 1976), pp. 67–137.

³In “Metaphysical Antirealism?” *Mind* 95 (1986), pp. 417–431; reprinted in *The Philosopher’s Annual* Vol. IX–1986 (Atascadero, Calif.: Ridgeview Publishing Co., 1988).

future.”⁴ What is it to “grasp” such a rule? What is it for me to have grasped a rule that determines that $68 + 57$ yields the answer 125, in the case that I have never happened to add 68 to 57? No such rule is determined merely by extrapolation from previous cases in which I have applied “+” to pairs of numbers; there are always infinitely many functions that accord with a given finite list of such argument-argument-value trios. For example, the “quus” rule might accord:

$$\begin{aligned} x \text{ quus } y &= x + y, \text{ if } x, y < 57 \\ &= 5 \text{ otherwise}^5 \end{aligned}$$

Nor (and this is more obviously a Wittgensteinian theme) can we suppose that my meaning addition by “+” consists in my having given myself general directions for what to do when encountering “+.” To give myself general directions would be to lay down a rule of procedure for myself. What then constitutes my meaning by this set of instructions, by this laid-down rule, one procedure rather than another? Certainly this set of instructions does not include a thought of each of the infinitely many sums there are. And my past performances when having this set of instructions in mind do not exemplify a unique general procedure but many such possible procedures. Supplementing the instructions with another set of instructions explaining how to follow the first set leads only to a regress. How then *is* the correct interpretation of the instructions in my mind determined?

Changing the example, Kripke writes

It has been supposed that all I need to do to determine my use of the word ‘green’ is to have an image, a sample, of green that I bring to mind whenever I apply the word in the future. When I use this to justify my application of ‘green’ to a new object, should not the sceptical problem be obvious to any reader of Goodman? Perhaps by ‘green’, in the past I meant *grue*, and the color image, which indeed was *grue*, was meant to direct me to apply the word ‘green’ to *grue* objects always. If the *blue* object before me now is *grue*, then it falls in the extension of ‘green’, as I meant it in the past. It is no help to

⁴S. Kripke, *Wittgenstein on Rules and Private Language* [hereafter *WORPL*] (Cambridge, Mass.: Harvard University Press, 1982), p. 7.

⁵*Ibid.*, p. 9.

suppose that in the past I stipulated that 'green' was to apply to all and only those things 'of the same color as' the sample. The sceptic can reinterpret 'same color' as same *schmolor*, where things have the same *schmolor* if . . .⁶

Now it is true that arguments of this sort take hold only if we reject the possibility that intentionality is a *sui generis* feature given to consciousness.⁷ We must assume that what comes before the mind, whatever it is that enters or informs consciousness when one means something, does not *itself* determine a use for itself, a purpose for itself, a particular kind of connection that it is to have with one's activities. Rather, whatever comes before the mind is, in this respect, not different from any other item standing alone: "And can't it be clearly seen here that it is absolutely inessential for the picture to exist in his imagination rather than as a drawing or model in front of him? . . ." ⁸ Wittgenstein argues against the possibility that intentionality is a *sui generis* feature, by showing, for each of a series of cases, that the results of introspection when one means, understands or is guided in accordance with rules, are not the only or the final criteria that we use to determine what we mean or when we understand or are being so guided. What lies before consciousness does not determine its own significance; knowing what one means is not a matter, merely, of apprehending the contents of one's mind. In short, meaning is neither a state of awareness nor an epistemological given. It does not occur encapsulated within consciousness; it is not a state that simply *shows* its content or its significance. If there *is* such a thing as meaning something, say, meaning addition, its nature must lie in part in what is *not* simply given to consciousness.

Nor is it merely because the object thought of or meant is external to mind that meaning has an ingredient not given to consciousness. Meaning to perform a mental activity like adding in the

⁶Ibid., p. 20. Kripke's ellipsis points at the end; Kripke's footnotes omitted.

⁷Loar claims that Kripke has not demonstrated that intentionality is not this. Kripke's text does however contain several footnotes commenting on the relevant arguments in Wittgenstein's text. I mention these arguments below.

⁸L. Wittgenstein, *Philosophical Investigations* (New York, N.Y.: The Macmillan Company, 1953), Paragraph 141.

head, that is, having intentions about one's own thoughts, is fully infected with this non-given ingredient. Thus the problem posed is no different for the purest idealist than for the metaphysical realist. Nor is it only "naturalistically inclined philosophers" who need a solution to the Kripke-Wittgenstein paradox. It is anyone who has been convinced by Wittgenstein to doubt Brentano—or, say, convinced after Sellars to reject epistemological "givenness" in *all* of its multifarious forms.

Could it be that the non-given ingredient that pins down what rule I intend to follow for "+" is the *disposition* I have to proceed in a certain way when encountering "+"? Setting aside the problem of what Wittgenstein may have intended as an answer to this question, surely Kripke is right to answer no. Kripke gives two main reasons for his answer. First, people are in fact disposed to make mistakes in arithmetic. Second, the addition function applies to numbers of any magnitude, but "some pairs of numbers are simply too large for my mind—or my brain—to grasp."⁹ Nor will it help to take into account dispositions I may have to correct myself or to accept correction from others. Some of my dispositions are dispositions to miscorrect myself. (I often do this when trying to add long columns of figures.) And there are surely conditions under which I would be disposed to accept miscorrection from others.

Kripke concludes, or he claims that Wittgenstein concludes, that there is, indeed, *no fact to the matter of what I mean by "+"*. This conclusion is what I am calling the "Kripke-Wittgenstein paradox."¹⁰ Wittgenstein, Kripke claims, offers only a "sceptical solu-

⁹WORPL, pp. 26–27.

¹⁰Kripke places a great deal of emphasis on the failure to find anything that "justifies" my proceeding as I do when I follow a rule, and he seems to think of a "justification" as something that must be, by its very nature, open to or within consciousness. Similarly: "Even now as I write, I feel confident that there is something *in my mind* [italics mine]—the meaning I attach to the 'plus' sign—that *instructs* me [italics Kripke's] what I ought to do in all future cases" (WORPL, pp. 21–22). And "The idea that we lack 'direct' access to the facts whether we mean plus or quus is bizarre in any case. Do I not know, directly, and with a fair degree of certainty, that I mean plus?" (WORPL, p. 40). Indeed, many passages in Kripke's essay suggest that what bothers him the most is not that nothing seems to determine what rule I am following, but that nothing *before my mind* determines it. The feeling is conveyed that Kripke finds the real blow to be that the intentionality involved in rule following does not reside *within* conscious-

tion” to this paradox, a solution that “begins . . . by conceding that the sceptic’s negative assertions are unanswerable.”¹¹ I propose to offer a “straight solution” to this paradox, one that “shows that on closer examination the scepticism proves to be unwarranted.”¹²

Kripke distills the essence of the failure of dispositional accounts to capture the nature of rule following thus:

A candidate for what constitutes the state of my meaning one function, rather than another, by a given function sign, ought to be such that, whatever in fact I (am disposed to) do, there is a unique thing that I *should* do. Is not the dispositional view simply an equation of performance and correctness? Assuming determinism, even if I mean to denote *no* number theoretic function in particular by the sign ‘*’, then to the same extent as it is true for ‘+’, it is true here that for any two arguments *m* and *n*, there is a uniquely determined answer *p* that I would give. (I choose one at random, as we would normally say, but causally the answer is determined.) The difference between this case and the case of the ‘+’ function is that in the former case, but not in the latter, my uniquely determined answer can properly be called ‘right’ or ‘wrong’.¹³

.

The fundamental problem . . . is . . . : whether my actual dispositions are ‘right’ or not, is there anything that mandates what they *ought* to be?¹⁴

ness. If *that* is what Kripke takes to be the root “Wittgenstein paradox,” then all will agree that Wittgenstein made no attempt to give a “straight solution” to it. Nor will I. On the other hand, if that were the main paradox, no one would ever have supposed that a dispositional account would be a “straight solution” to it and Kripke’s discussion of dispositional accounts should have been placed not with his account of “Wittgenstein’s paradox” but with his discussion of “Wittgenstein’s sceptical solution,” that is, as an account of what Kripke believed this skeptical solution was not.

Margaret Gilbert suggests (in conversation) that one paradox may be that meaning strikes one as being something that can be fully constituted at a given time *t*, whether or not meaning is something that happens within consciousness. And, to be sure, dispositions are usually taken to exist at given times so that a dispositional account might be viewed as an attempt at a straight solution to Gilbert’s paradox. Gilbert’s paradox, if one finds it paradoxical, is another that I will not attempt to solve. I will merely try to show how there is a fact to the matter of what I mean by “+.”

¹¹WORPL, p. 66.

¹²Ibid.

¹³Ibid., p. 24; footnotes omitted.

¹⁴Ibid., p. 57.

The problem is to account for the *normative* element that is involved when one means to follow a rule, to account for there being a *standard* from which the facts, or one's dispositions, can diverge.

II. GENERAL FORM OF THE SOLUTION

In the case of meaning, the normative element seems to be the same as the purposive element: to mean to follow a certain rule is to have as a purpose to follow it. Whether my actual dispositions are "right" or "wrong" depends on whether they accord with what I have purposed. The possible divergence of fact from a standard is, in this case, simply the failure to achieve a purpose.

Now having as one's purpose to follow a rule might involve having a representation of that purpose in mind, for example, in one's language of thought. But as Wittgenstein observed, any such representation would itself stand in need of interpretation. It would stand in need of a prior rule governing how it was to be taken, that is, how it was to guide one. And that one was to follow this prior rule could not *also* be a represented purpose, not without inviting a regress. To understand what it is to have an explicit purpose that one represents to oneself we must first understand what it is to have a purpose the content of which is *not* represented. Basic or root purposes must be *unexpressed* purposes.

"Intend" strongly suggests an explicitly represented purpose, that is, a purpose that is thought about. So let me use the verb "to purpose" (with a voiced "s"; yes, it *is* in the dictionary) to include this more basic way of having a purpose. We can then put matters this way: root purposing is unexpressed purposing; our job is to discover in what this purposing consists. Let us also distinguish among three ways of conforming to a rule: (1) merely coinciding with a rule (this is the way in which we conform to "quus" rules and to rules to which we have mere dispositions to conform), (2) purposefully following an explicit or expressed rule, and (3) purposefully conforming to an implicit or unexpressed rule. Way (3) involves having an unexpressed purpose to follow a rule and *succeeding* in this purpose. It is the same as displaying a *competence* in conforming to the unexpressed rule or displaying an *ability* to conform to it. Another way to explain our task, then, is to say that we need to learn what a competence in conforming to an unexpressed rule consists in, and how it differs from a mere disposition to coincide with the rule.

My thesis will be that the unexpressed purposes that lie behind acts of explicit purposing are biological purposes; a competence to conform to an unexpressed rule is a biological competence. By a biological purpose I mean the sort of purpose the heart has, or those of the eyeblink reflex, and the human brain. The purposes of these are functions that they have historically performed which have accounted for their continued proliferation. Biological purposes are, roughly, functions fulfilled in accordance with evolutionary design. It does not follow that capacities to perform biological functions are, in general, innate. For example, it is surely in accordance with evolutionary design that the newly hatched chick follows its mother about, but the chick is not born with that disposition. It is not born knowing which *is* its mother, but must imprint on her first. Yet the imprinting, and hence the following, both take place in accordance with evolutionary design. (Later in this paper I will devote considerable space to clarifying how even quite novel biological purposes can emerge as a result of experience and learning.)¹⁵

Suppose that explicit intending involves something like representing, imaging, or saying something to oneself and then using, or reacting to, or being guided by this representation in a purposeful way, that is, in a way that expresses a competence. My thesis, then, is that the purpose that informs this reacting, that makes it into a competence, is a biological purpose. Similarly, if knowing a language involves having a competence in following certain rules for construction and interpretation of sentences, the purpose that informs this competence, I will argue, is a biological purpose.

III. PURPOSIVE RULE FOLLOWING; COMPETENCE TO FOLLOW A RULE

Let me begin with a very simple example of an organism that displays a competence in conforming to a rule. According to the biologists Collett and Land,

¹⁵Full details of the notion of biological function that I rely on in this paper are given under the label "proper functions" in my *Language, Thought, and Other Biological Categories* [hereafter *LTOBC*] (Cambridge, Mass.: Bradford Books/The MIT Press, 1984), Chapters 1–2. See also my "In Defense of Proper Functions," *Philosophy of Science* 56 (1989), pp. 288–302.

TRUTH RULES AND HOVERFLIES

Males of many species of hoverfly spend much of the day hovering in one spot, thus keeping their flight muscles warm and primed so that they are ready to dart instantly after any passing female that they sight. This chasing behavior is on such a hair-trigger that all manner of inappropriate targets elicit pursuit (pebbles, distant birds, and midges so small as to be scarcely visible to a human observer) as well as a very occasional female. Although selective pressures have favored a speedy response above careful evaluation of the suitability of the target . . . the response itself is precisely tailored to optimize the capture of objects which are roughly the same size and speed as a conspecific.¹⁶

Rather than turning toward the target in order to track it, the hoverfly turns away from the target and accelerates in a straight line so as to intercept it. Given that (1) female hoverflies are of uniform size, hence are first detected at a roughly uniform distance (about .7 m), (2) females cruise at a standard velocity (about 8m/sec), and (3) males accelerate at a constant rate (about 30–35m/sec²), the geometry of motion dictates that to intercept the female the male must make a turn that is 180° away from the target minus about 1/10 of the vector angular velocity (measured in degrees per second) of the target's image across his retina. The turn that his *body* must make, given as a function of the angle off center of the target's image on his retina, equals the (signed) angle of the image minus 1/10 its vector angular velocity, plus or minus 180°. According to Collett and Land, whether it is dried peas, male hoverflies, female hoverflies or flying blocks of wood that he spots, that is exactly the rule to which the hoverfly conforms. Taking note that this rule is not about how the hoverfly should behave in relation to distal objects, but rather about how he should react to a proximal stimulus, to a moving spot on his retina, let us call this rule "the proximal hoverfly rule."

I have chosen the proximal hoverfly rule as my first example of rule following because it seems so unlikely that the hoverfly calculates over any inner representation of this rule in order to follow it. Rather, the hoverfly has an unexpressed biological purpose to conform to this rule. That is, the hoverfly has within him a genetically determined mechanism of a kind that historically proliferated

¹⁶"How Hoverflies Compute Interception Courses," *Journal of Comparative Physiology* 125 (1978), pp. 191–204.

in part *because* it was responsible for producing conformity to the proximal hoverfly rule, hence for getting male and female hoverflies together. This mechanism may account for various other dispositions of the hoverfly, for example, causing him to attract predators by his conspicuous darting movements, or causing characteristic uniform mathematically describable patterns to play on his retina as he turns after the female. But mentioning these latter dispositions does not help to explain why the mechanism has survived, why it has proliferated in the species. Conformity to the proximal hoverfly rule, on the other hand, has helped to explain the reproductive success of (virtually) every ancestor hoverfly, hence to explain the continued presence of the mechanism in the species. Conformity to the proximal hoverfly rule, then—not attracting predators or producing certain patterns on the retina—is a biological purpose of this mechanism, hence of the hoverfly. For similar reasons, a biological function of the heart is to pump blood but not also, say, to make a jazzy sound, and a biological function of the eyeblink reflex is to cover the eyes momentarily, but not also to swing the eyelashes in a graceful arc away from entanglement with the eyebrows, nor to point with them at the navel.

The hoverfly displays a *competence* in conforming to the proximal hoverfly rule when his coinciding with it has a “normal explanation,” that is, an explanation that accords with the historical norm. That his behavior coincides with the rule must be explained in the same way, or must fit the same explanation schema, that accounted in the bulk of cases for the historic successes of his ancestors in conforming to the rule. Presumably this normal explanation makes reference to the way the hoverfly’s nervous system is put together, how it works, how it is hooked to his retina and muscles, etc. If the hoverfly ends up coinciding with the rule not because his nerves and muscles work in a normal way but only because the wind serendipitously blows him around to face the right direction, he fails to express a competence.¹⁷

Not just anything a human effects is a human action. Effects that are actions must be intended, or at least foreseen, and must be

¹⁷For a full discussion of normal (“Normal”) explanations for performance of proper functions, see *LTOBC*, Chapters 1 and 2, and my “Bio-semantics,” *The Journal of Philosophy* 86 (1989), pp. 281–297.

generated from intentions in a normal way. Effects of human bodily movements that are not actions are called “accidental.” Similarly, not just any process that originates in an animal’s organs or behavioral systems is a biological activity. Biological activities are only those that express competencies. They correspond to normally fulfilled biological purposes, that is, to what the animal does in accordance with evolutionary design. Conversely, behavior that fails to express a competence corresponds to what an animal effects, biologically, only by accident. Thus the heart’s saying pit-a-pat, the eyelashes’ moving away from the eyebrows in a graceful arc and the hoverfly’s coinciding, but due only to the wind, with the proximal hoverfly rule are not biological activities, but biological accidents.

To say that a given male hoverfly has a biological purpose to conform to the proximal hoverfly rule is very different from saying either that he himself has a history of having conformed to it (perhaps he has just reached adolescence) or that he has a disposition to conform to it. The normal hoverfly has a disposition to dart off when it sees a flying bird—and also a disposition to squash when stepped on—but these dispositions do not correspond to biological purposes or to competencies. Conversely, male hoverflies that are crippled or blind have no disposition to conform to the proximal hoverfly rule, but still it is one of their biological purposes to do so. As male members of the hoverfly species, conforming is the biological norm, the standard for them.¹⁸

To say that the hoverfly has as a biological purpose to follow the proximal hoverfly rule is also quite different from saying that this rule is the only rule that fits all past instances of hoverfly turns, say, that resulted in hoverfly procreation. Suppose it were so that never in history had a male hoverfly spotted a female that happened to approach him at such an angle as to produce an image on his retina with a clockwise angular velocity between 500° and 510° per second. Then the proximal *quoverfly* rule, “If the vector angular velocity of the target’s image is *not* counterclockwise and between 500° and 510° per second, make a turn that equals the (signed) angle of the image minus $1/10$ its vector angular velocity,

¹⁸On the proper functions of imperfect members of a biological category, see *LTOBC*, Chapters 1 and 2.

plus or minus 180°; at ease otherwise,” fits all past actual cases of successful female encounters. But it is not a rule the hoverfly has as a biological purpose to follow. For it is not because their behavior coincided with *that* rule that the hoverfly’s ancestors managed to catch females, hence to proliferate. In saying that, I don’t have any particular theory of the nature of explanation up my sleeve. But surely, on any reasonable account, a complexity that can simply be dropped from the explanans without affecting the tightness of the relation of explanans to explanandum is not a *functioning* part of the explanation. For example, my coat does not keep me warm because it is fur-lined *and red*, nor because it is fur-lined *in the winter*, but just because it is fur-lined. (True, I am making the assumption that the qualifications and additions that convert the proximal hoverfly rule into the proximal quoverfly rule are objectively qualifications and additions rather than simplifications. This assumption rests upon a metaphysical distinction between natural properties and kinds and artificially synthesized grue-like properties and kinds or, what is perhaps the same, depends upon there being a difference between natural law and mere *de facto* regularity. But my project is to solve the Kripke-Wittgenstein paradox, not to defend common-sense ontology. Nor should either of these projects be confused with solving Goodman’s paradox.)¹⁹

To say that the hoverfly has as a biological purpose to follow the proximal hoverfly rule is also quite different from saying that this rule is the only rule that fits the actual dispositions of normal hoverflies or of past hoverflies that managed to procreate. Suppose that, given the principles in accordance with which the hoverfly’s turn-angle-determining devices work, engineering constraints necessitated a mechanism normal for hoverflies with a blind spot for clockwise angular velocities between 500° and 510° per second. These particular velocities produce no reaction at all on the part of the male. Then the same proximal quoverfly rule

¹⁹Goodman’s paradox is a paradox in epistemology. Kripke, on the other hand, is concerned not about how we could know or discover what someone means by “plus” but about what this determinate meaning *consists* in. Note too that assuming common-sense ontology does nothing, by itself, toward solving Goodman’s paradox, which concerns how we can *know* or reasonably guess which entities are the basic ontological ones, *supposing* there to be such.

mentioned above fits the actual dispositions of all normal hoverflies, but it still would not be a rule that the hoverfly has as a biological purpose to follow. The hoverfly's biological purposes include the expression only of dispositions that have helped to account for the proliferation of his ancestors. By hypothesis, the disposition to rest at ease when the target's image is counterclockwise and between 500° and 510° per second did not help the hoverfly's ancestors to propagate. It was only the times that the proximal hoverfly rule was obeyed that the ancestors procreated. So the hoverfly resting at ease behind his blind spot is not displaying a competence. It is conformity to the proximal hoverfly rule, not the quoverfly rule, that he biologically purposes, even if normal hoverflies are not especially accurate in fulfilling this natural purpose, in conforming to this ideal.

IV. PROXIMAL VS. DISTAL RULES

My plan, as I have indicated, is slowly to make plausible the claim that the normative element that is involved when one means to follow a rule is biological purposiveness. Meaning to follow a rule differs from having a disposition to coincide with a rule, in the same way that the hoverfly's biologically purposing to follow the proximal hoverfly rule differs from having a disposition to coincide with it. That is how I aim to solve the Kripke-Wittgenstein paradox concerning what *constitutes* rule following. At the same time, however, I wish to build a case that language understanding or language competency is competency in the biological sense. And I wish to argue that it is possible to have a biological competence to follow correspondence truth rules, hence that a "realist" theory of language understanding is possible on the biological model. To gain this latter end, we need to discuss distal as well as proximal rules.

Conforming to the proximal hoverfly rule is a means, for the hoverfly, of following a less proximal, or more distal rule: "If you see a female, catch it." Call this "the distal hoverfly rule." To say that conformity to the proximal hoverfly rule is a means to conformity to the distal rule is the same as to say that the mechanism that has historically accounted for the overwhelming majority of ancestor hoverflies' successes at conformity to the distal hoverfly rule begins with conformity to the proximal rule. That is, the normal explanation for conformity to the distal rule contains

the specification that the hoverfly first conform to the proximal rule.

Now whether the hoverfly succeeds in following the proximal hoverfly rule depends, for the most part, only upon whether his insides are working right, that is, on whether he is a normal healthy member of his species. But whether, or how often, he manages to conform to the distal hoverfly rule depends upon more. It depends upon conditions that are outside his body and over which he has no control, such as how hard the wind is blowing, whether the females that pass by are in fact of normal size, traveling at the normal speed and, perhaps, whether they are willing. Without doubt, then, hoverflies are worse at conforming to the distal than to the proximal hoverfly rule. That is, their competence or ability to conform to the distal rule is less reliable than their competence or ability to conform to the proximal rule. But that the hoverfly may not be very reliable in his conformity to the distal hoverfly rule bears not at all upon whether it is one of his biological purposes to conform. Compare: it is a biological purpose of the sperm to swim until it reaches an ovum. That is what it has a tail for. But very few sperm actually achieve this biological end because ova are in such short supply. Reaching an ovum is a purpose of the sperm since it is only because ancestor sperm reached ova that they reproduced, thus proliferating the tail. Similarly, it was only when ancestor hoverflies conformed to the distal hoverfly rule that they *became* ancestors.

Turning the coin over, the hoverfly is very reliable in his coincidence with this “overkill rule”: “Dart off after everything that flies by you subtending about $.5^\circ$ on your retina, whether it’s male, female, animate or inanimate, bird, plane or Superman.” But this overkill rule does not correspond to any biological purpose of the hoverfly. True, conforming to the proximal hoverfly rule is one of the hoverfly’s biological purposes, and conforming to this rule will *result* in his coinciding with the overkill rule if there are objects other than female hoverflies flying about him (even if there are not). But it is not coinciding with the overkill rule that has helped to account for hoverfly proliferation. Only the times when the distal hoverfly rule was obeyed did hoverfly ancestors procreate.²⁰

²⁰More precisely, only the distal hoverfly rule would be mentioned in giving a “most proximate normal explanation” of the function of the

It is conformity to the distal hoverfly rule that explains the ancestor hoverflies' *successes*. As the hoverfly chases after a distant bird, he expresses no competence except, of course, competence to conform to the proximal hoverfly rule. Conformity to the distal hoverfly rule, not to the overkill rule, is what he biologically purposes, though at the moment he is accidentally, that is, nonbiologically, doing something else.

That is how purposes inform the rule-following behavior of the hoverfly, how norms, standards, or ideals apply to his behaviors, hence how the hoverfly comes to display competences or abilities to conform to rules rather than mere dispositions to coincide with them.²¹ But the unexpressed rules that humans purposively conform to, at least most of those that they purposively conform to when using inner or outer language, are not rules that they are genetically hard-wired to follow, but rules that they have learned. How then can humans biologically purpose to follow such rules? Before turning directly to the problem of human rule following, let us examine a more simple case of learned biological purposes, of learned competence—the case of a simpler animal that learns to follow rules.

V. LEARNED OR DERIVED RULES AND COMPETENCES

If a rat becomes ill within a few hours after eating a specific food, it will later shun all foods that taste the same. For example, if the rat eats soap and soon becomes ill, thereafter it will refuse to eat soap. Although the rat may have dragged certain nesting materials home or explored new territory just before becoming ill, it will not on that account shun that kind of nesting material or that territory. Nor will it shun foods that merely look the same or that are found in the same place as the food eaten prior to illness.²² It

hoverfly's turning mechanism. See *LTOBC*, and the discussion of "normal explanations" in my "Biosemantics."

²¹Notice that it is the reference to evolutionary *history* that has been doing all of the work in explaining how norms come to apply to the activities of an animal, in explaining how there can be a standard from which the facts of individual behavior diverge. I defend the position that function always derives from history in "In Defense of Proper Functions."

²²The reference is to studies by John Garcia. A bibliography of his papers may be found in *The American Psychologist* 35 (1980), pp. 41–43.

thus appears that a quite specific mechanism is harbored in the rat, a proper function of which is to produce conformity to the specialized rule "If ingestion of a substance is followed by illness, do not ingest any substance with that taste again." Call this rule the "proximal rat rule." Clearly, following the proximal rat rule is a biological means to following a more distal rat rule, say, "Do not eat poisonous substances"; helping to produce conformity to this rule is a further proper function of the relevant inborn mechanisms in the rat.

Now the proximal rat rule, like the proximal hoverfly rule, tells the animal what to do given certain experiential contingencies. There is a difference, however, in the normal manner of executing these two rules. When the hoverfly conforms to his rule, nothing in his body undergoes a permanent change, but this is not so in the case of the rat. Suppose, for example, that the rat has just become ill after eating soap. In order to conform to the proximal rat rule, in order to avoid henceforth what tastes like soap, the rat's nervous system must first conform to certain preliminary "rules," rules that dictate that a certain sort of permanent change take place in it. The rat, we say, must "learn" in order to conform to his rule. But the fact that the rat's evolutionary history dictates that it is normal for him to undergo learning in order to follow his rule rather than following it directly does not affect the biological status of the rule. That he should follow his rule is one of his biological purposes for exactly the same reason that the hoverfly's rule following is biologically purposed. Conformity to the rat rule is what his ancestor rats had in common in those cases in which possession of the relevant inborn mechanisms aided them to flourish and proliferate, so it is what the mechanism, hence the rat, biologically purposes.

Now the rat that conforms to the proximal rat rule, if he ever becomes ill after eating, ends by conforming to a *derived* proximal rat rule, say, the rule "Do not eat what tastes like soap." Indeed, if a rat becomes ill after eating soap, it immediately becomes one of his biological purposes to follow the rule "Do not eat what tastes like soap." For that he is to follow this derived rule is logically entailed by the proximal rat rule plus the premise that he has in fact become ill after eating soap. Similarly, the hoverfly that currently has an image of appropriate size traversing his retina at a 60° angle with an angular velocity of 100° per second currently has as a biological purpose to make a turn of 130° . Notice that the hoverfly

has this biological purpose quite independently of whether or not any hoverfly has ever been in exactly this experiential position before. It is theoretically possible, even if unlikely, that no hoverfly has ever had exactly *this* biological purpose before. This is similar to our rat who is sick after eating soap. It is now one of his biological purposes to follow the derived proximal rat rule "Do not eat what tastes like soap" even if it should be true that no other rat in history has ever become sick after eating soap, hence true that no rat in history has ever had this particular biological purpose before.

In this manner, animals that learn can acquire biological purposes that are peculiar to them as individuals, tailored to their own peculiar circumstances or peculiar histories.²³ Although biological purposes are functions fulfilled in accordance with evolutionary design, they need not be innately given purposes. Similarly, biological competences need not be innate. A proper biological activity of an animal can be something that experience has prompted or "taught" the animal to do, experience coupled with an innate mechanism for being guided to learn by experience.

Nor is there need for such mechanisms to be as specialized as the mechanism that conforms the rat to the proximal rat rule. Not long ago many learning theorists believed that all animal learning took place in accordance with principles that were not species specific but universal. Suppose that this were true. Suppose that every species learned in accordance with the principles of one person's favorite general theory of operant conditioning, so that no reference to the particular evolutionary niche of a species was ever needed to explain how its learning mechanisms had historically enhanced its fitness. Then there would have to be some rarefied hypergeneral explanation of how and why these learning principles worked. Such an explanation might make reference, for example, to specific principles of generalization and discrimination used in differentiating stimuli and in projecting what is to count as "the same" behavior again, that is, reference to universal proximal rules followed during learning. It would have to tell how

²³A much more detailed discussion of "derived proper functions" may be found in *LTOBC*, Chapter 2, and in my "Thoughts Without Laws: Cognitive Science With Content," *The Philosophical Review* 95 (1986), pp. 47–80.

and why these particular ways of generalizing and discriminating effected, often enough, isolation or zeroing in on sufficiently reliable causes of reinforcement, and in what universal manner (!) reinforcers are connected with the well-being of animals. Thus it would tell *how* possession of the universal mechanism had normally, that is, historically, enhanced fitness in animals generally. Specific applications of this general explanation schema to individual animals in individual circumstances would then determine which among the various effects of their motions were the proximal and distal biological activities of these individuals, as they learned and applied their learning. Such applications would determine, for example, what specific reliable causes of reinforcement were purposefully being zeroed in on by particular animals at particular times, that is, what these animals were “trying” to learn and, after they learned it, what the specific goals of their learned behaviors were.

Now it is important to note that to fulfill a biological purpose is not always to take a step towards flourishing or propagating; it is not always good for an animal to fulfill its biologically determined goals. For example, a rat might come to have as a biological purpose to follow the derived rule “Do not eat what tastes like soap” even if it were true (I suspect it is true) that soap does not *make* rats sick or does not poison them. Suppose, rather, that the rat eats soap and then becomes ill due to a bout with *Rattus enteritis*. Still, in order to conform to the proximal rat rule, he must now conform to the derived proximal rule “Do not eat what tastes like soap,” for this derived rule is entailed by the proximal rat rule given his situational experience. Yet following this derived rule may, in fact, have no tendency to bring him into conformity with the more distal rat rule “Do not eat poisonous substances.” So it can happen that the rat acquires a biological purpose and acquires a competence to conform to a derived rule which does not further the end that is this rule’s own *raison d’être*. Indeed, the rat *could* acquire a derived purpose and a competence to behave in a manner that was actually detrimental to him, say, a competence to follow the rule “Do not eat what tastes like corn” when, in fact, unless he eats corn, given his circumstances, he will starve. Compare: the hoverfly, dutifully conforming to the proximal hoverfly rule (the rule that tells how he is to react to a moving image on his retina) may thereby dart off after a bird, who would not otherwise

have spotted the hoverfly, hence would not have *eaten* him. Thus it is that an individual may have a biological purpose and a competence to follow a derived rule that has no tendency to further the interests either of the individual or of his species and, more specifically, no tendency to produce conformity to more distal rules toward which following it was, biologically, supposed to be a means.²⁴

What an animal is doing in accordance with evolutionary design need not be anything that any member of its species has ever done before. And it need not be anything that is good for the animal to do. So surely it need not be anything that common sense would call “natural” for it to do. Consider a circus poodle riding a bicycle. It is performing what common sense would call a most “unnatural” act. Yet it is one of the dog’s biological purposes to perform that act. Biologically, the (typical circus) dog’s distal action is procurement of his dinner. The dog harbors within him an intricate mechanism, operating in accordance with certain largely unknown but surely quite definite and detailed principles, in accordance with which dogs have been designed to develop perceptual, cognitive and motor skills and to integrate them so as to effect procurement of dinner in their individual environments. Living in an unusual environment, the circus dog acquires unusual purposes and competences when he applies his “dog rules” to his environment. But, although he may be making the audience laugh by accident, he is certainly not balancing on that bicycle by accident. He is balancing purposefully or in accordance with evolutionary design—in accordance with another application of the same general principles that procured his ancestors’ dinners during evolutionary history.

VI. HUMAN RULE FOLLOWING

Humans are very sophisticated creatures, so we tell ourselves. We not only learn but learn new ways to learn, develop new concepts, and so forth. Further, much of our behavior results not just from learning but from theoretical and practical inference. But there must still be a finite number of inborn mechanisms, oper-

²⁴For further details on conflicting proper functions, see *LTOBC*, Chapter 2, and my “Thoughts Without Laws.”

ating in accordance with a finite number of natural principles, having a finite number of biologically proper functions, that account for our dispositions to do these things. Coordinately, there must be a finite number of proximal and distal “*Homo sapiens* rules” that we have as biological purposes to follow, and there must be mechanisms to implement these rules built into the basic body and brain of normal persons.

Consider then any bit of human behavior produced by biologically well-functioning behavior-regulating systems, by systems that are not broken or jammed. (Behavior that results from malfunction is, of course, overwhelmingly unlikely to bear fruit of any interesting kind.) There will be a way of describing this behavior that captures its aspect as a *biological* activity, a description that tells what proximal and distal biological purposes, and what biological competences if any, the behavior expresses. This will be so even if the behavior is totally unique, or systematically self-destructive, or not “natural” by any common-sense standards. But of course there will also be numerous ways of describing the behavior that fail to express its biological purposiveness, many “quus-like” descriptions — as “pointing toward the navel with the eyelashes” quus-describes the eyeblink reflex. So the question arises, what is the relation of *ordinary* human purposes, of human intentions and meanings, to biological purposes? Are descriptions of human intentional actions quus-descriptions from the standpoint of evolutionary design? Do ordinary human intentions merely, accidentally, *cohabit* with biological purposes?

Surely a naturalist must answer no. Ordinary human purposes, ordinary intentions, can only be a *species* of biological purpose. To suppose otherwise would be to suppose that the whole mechanism of human belief, desire, inference, concept formation, etc., the function of which culminates in the formation and execution of human intentions, is, as functioning in this capacity, an epiphenomenon of biology, an accidental by-product of systems that nature designed for other purposes. And what would these other purposes be?²⁵

This accords with conclusions we reached earlier on the nature of explicit intentions. Explicitly meaning or intending, if this re-

²⁵For a more detailed defense of this claim, see my “Biosemantics.”

quires representing what one intends, presupposes a prior purposing: purposing to let the representation guide one in a certain way. This is true whether we are talking about representation in an inner medium, say, in a “language of thought,” or representation in a public medium—talking, say, about the use of “plus.” But this prior purposing cannot be analyzed as the original explicit purposing was analyzed without regress. Rather, a prior unexpressed purposing must be assumed. The reasonable conclusion seems to be that ordinary explicit intending rests on biological purposing—biologically purposing to be guided by, to react this way rather than that to, one’s representations. Whether this biological purposing is innate (compare Fodor’s version of the “language of thought”) or whether it is derived via learning, mechanisms of concept formation, etc., it must *ultimately* derive its content from the details of our evolutionary history.

So unless doing arithmetic results from a total breakdown of the cognitive systems (in which case there may be nothing you purpose when you encounter “plus”: how you react to it is accidental under every description) then *whatever* you mean to do when you encounter “plus,” that content has been determined by your experience coupled with evolutionary design. But, reasonably, whatever you mean by “plus” is the same as what other people mean who are endowed with the same general sort of cognitive equipment and have been exposed to the same sort of training in arithmetic. This meaning has been determined by the application of *Homo sapiens* rules of some kind to experience. It is likely that these are extremely abstract general purpose *Homo sapiens* rules, in accordance with which human concept formation takes place, and it is likely that the explanation of the efficacy of these rules makes reference to very deep and general principles of ontology. But it is not my task to speculate about the precise form these *Homo sapiens* rules take, or about how the experience of standard training in arithmetic elicits from them the capacity to mean plus. Speculation about the specific forms that our most fundamental cognitive capacities take is the psychologist’s job.²⁶

²⁶But, people still persist in asking, How do you know that we really *do* end up meaning *plus* by “plus”? How do you know we don’t mean *quus*? Because if we meant *quus* then “plus” would mean *quus*, and the way to *say* that we all meant *quus* would be “we all mean plus”—which is what I

I believe that these considerations constitute, albeit in very rough and broad outline, the solution to the Kripke-Wittgenstein paradox.

VII. TRUTH RULES: VERIFICATIONIST OR CORRESPONDENCE?

I have sketched a theory about meaning in the sense of purposing—both expressed and unexpressed purposing. It remains to connect this theory with the theory of semantic meaning.

Truth rules are rules that project, from the parts and structure of sentences in a language, the conditions under which these sentences would be true. Such rules express, of course, an aspect of the meaning of the sentences. The question is whether the conditions referred to by truth rules are to be understood in a “realist” way as correspondence truth conditions, or in a “verificationist” way as assertability conditions. Dummett’s concern about truth rules is this: whatever connection there is between sentences and that which determines their truth has to be a connection that is established via the actual employment of the language. Whatever form truth rules take, realist or verificationist, the *practical* abilities of speakers who understand a language must reflect these rules, indeed, must determine their content. Hence an analysis of the structure of the abilities required for language use and understanding should reveal the kind of rules truth rules are. But, Dummett argues, the only truth rules we could possibly exhibit a practical grasp of are verificationist truth rules.

In Section VIII below, I will claim that Dummett’s argument hangs on treating language abilities or competencies, hence the following of language rules, as mere dispositions, or alternatively (perhaps), as taking place wholly within consciousness, and I will add to the arguments already piled up by Wittgenstein and Kripke against the adequacy of this sort of treatment. In the present section, however, I wish to propose a positive thesis. My claim will be that if we interpret rule following and, in general, purposes and competencies in the biological way, then we can see how, on the

said. Compare Donald Davidson, “Knowing One’s Own Mind,” *Proceedings and Addresses of the American Philosophical Association* 60 (1987), pp. 441–458, and Tyler Burge, “Individualism and Self Knowledge,” *Journal of Philosophy* 60 (1988), pp. 649–663.

contrary, reference to correspondence truth rules might *easily* fall out of an analysis of language competence.

We begin by observing that whatever the content of truth rules may be, realist or verificationist, the intent or purpose of anyone engaged in making sincere assertions in a language must be to conform their sentences to these rules. The sincere speaker purposes to make assertions that are true. It follows that the *way* that the actual practice of a language embodies truth rules is that these are the rules in accordance with which the competent speaker (or thinker), when sincere, purposes to make (or think) assertions. These are rules that he is, as it were, *trying* to follow insofar as he is sincerely speaking (or thinking) *that* language. On the bottom layer at least (perhaps the layer that governs the language of thought) these rules must of course be unexpressed rules. But precisely because truth rules are at bottom unexpressed rules, introspection can give us no handle on what kind of rules they are. Rather, it is necessary to develop a *theory* about truth rules, an explanatory hypothesis about what rules we are purposing to follow when we make sincere assertions.

Assuming a biological standpoint, the question whether truth rules are realist or verificationist can be expressed by asking how “proximal” vs. “distal” truth rules are. The proximal hoverfly rule was a rule about how the hoverfly was to respond to a moving image on his retina, that is, roughly, to sensory stimulations. The distal hoverfly rule was a rule about how the hoverfly was to end up interacting with his more removed environment, namely, with females that entered his life at a distance. “Verificationist” truth rules, as Putnam and Dummett envision these, would be rules that governed responses to prior thoughts and, as Dummett has put it, “bare sense experiences,” hence would be proximal rules.²⁷ “Realist” or correspondence rules, on the other hand, would for the most part be distal rules, rules that governed the manner in which assertions were to correspond to affairs that lie, very often, well beyond the interface of body and world. Convinced by Wittgenstein and Kripke that purposing to follow a rule is not something encapsulated *within* consciousness, we are not compelled to suppose that truth rules have to be rules about what is to happen either in the mind or at the interface between mind, or

²⁷“What is a Theory of Meaning? (II),” p. 111.

body, and world. So let us ask what it would be like if truth rules were distal correspondence rules.

The first thing to note is that if truth rules were distal rules they would surely have to be *backed* by proximal rules, rules about how to respond to our thoughts (inference) and to the immediate fruits of our perceptual explorations (perceptual judgment). They would have to be *backed* by rules that determined assertability conditions, the innermost of these conditions being within the mind or brain or at the interface of mind or brain and world. Call these back-up rules “proximal assertability rules.” Proximal assertability rules would concern the most proximal conditions under which we should say or think certain things. Conformity to these rules would have, as a biological purpose, to effect conformity to distal rules, that is, to correspondence truth rules. These truth rules would concern distal conditions under which we should say or think certain things. The truth rules might imply directives with this sort of form: if you have reason to speak (think) about the weather in Atlanta, say (think) “It is snowing in Atlanta” when and only when it is snowing in Atlanta; if you have reason to speak (think) about the color of snow, say (think) “Snow is white” if and only if snow is white. For a simple biological model here, compare worker honeybees. They (biologically) purpose to follow rules of this kind: when dancing, angle the axis of your dance 10° off the vertical if and only if there is a good supply of nectar 10° off a direct line from hive to sun. (Proposals concerning how humans might *learn* how to (purpose to) conform to distal correspondence truth rules are detailed in *LTOBC*.)²⁸

Conforming to the proximal hoverfly rule and the proximal rat rule often fails to bring hoverflies and rats into conformity to the distal hoverfly and rat rules. Similarly, conforming to proximal assertability rules might often fail to bring humans into conformity to truth rules. One can unknowingly say what is false even though one has good evidence for what one says. And one frequently fails to say what is true, indeed, to say anything at all, because one lacks any evidence at all, either for or against. Also, whether conformity to the proximal hoverfly and rat rules helps to produce conformity to the distal hoverfly and rat rules on this or that occasion

²⁸Chapters 9, 17 and 18.

often depends upon factors in the hoverfly's or rat's external environment over which it has no control. Similarly, whether conformity to proximal assertability rules would bring us into conformity to truth rules in this case or that might depend upon factors over which we had no control. For example, circumstances responsible for most perceptual illusions are circumstances outside the observer which, normally, he neither controls nor needs to control. Nor is not having enough evidence either to affirm or to deny a proposition typically something that it is within one's control to remedy. The principles in accordance with which biological devices perform functions that are proper to them always refer, in the end, to conditions external to these devices. These are conditions that have *historically* been present often enough to enable a critical proportion of ancestors of those devices to perform these functions, or to perform them a critical proportion of the time, but that cannot be counted on always to be present. All biological devices are fallible devices, even when normal and healthy.²⁹

It follows that the proximal assertability rules for a sentence would not *define* its semantics, for they would not determine what its truth conditions were. Rather, following proximal assertability rules would be means that were, merely, approximations to the end that was following correspondence truth rules—more or less helpful and more or less reliable means to that end. Let us reflect for a moment upon certain consequences of this model.

If proximal assertability rules were rules that we followed only as a more or less reliable means to following distal truth rules, then it would not at least be obvious that those who shared a language in the sense of having competences to abide by the same truth rules would have any need to share proximal assertability rules as well. The male hoverfly follows the distal rule "If you see a female, catch it" by following the proximal hoverfly rule. The male housefly follows the same distal rule by tracking the female rather than by plotting an interception path, employing different proximal means to the same distal end. Now consider how many different ways there are to make a map of a city: for example, by walking about with a yardstick, paper and pencil, by working from

²⁹For amplification of this very crucial theme, see my "Thoughts Without Laws" and "Biosemantics."

aerial photographs, by using surveyors' instruments, etc. Might there not also be various ways to make sentences that map onto the world in accordance with the same truth rules? Is there really any reason to suppose that only one set of proximal assertability rules could effect a reasonably reliable competence to conform to a given set of distal correspondence rules? Consider, for example, how many ways there are to tell whether a solution is acid or whether it has iodine in it. Consider how many alternative visual and tactile clues we use, on one occasion or another, to perceive depth. And consider: were the proximal assertability rules that Helen Keller used when she spoke English the same as those that you use? If not, does it follow that she did not really speak English after all?

Indeed, there is a sense, there is a way of individuating rules, in which it is impossible for people to share proximal assertability rules. Proximal assertability rules that I conform to correlate happenings at the periphery of my nervous system or body with sentences. Proximal assertability rules that you conform to correlate happenings at the periphery of your nervous system or body with sentences. For us to "share a set of proximal assertability rules" could not, of course, be for me to purpose to correlate happenings at the periphery of *your* body with *my* sentences. If I purposed to do that, I would be purposing to conform to a distal rule, not a proximal rule. We could "share proximal assertability rules" only in the sense that our rules ran parallel. But it is not immediately obvious what the point of running parallel to one another with language might be. Why would you take any interest in the sentences I uttered, if these correlated only with what was happening at the ends of my afferent nerves? Only if the proximal assertability rules that you and I used effected relatively reliable conformity to the same *distal* correspondence rules would there be any point in talking to one another. But if agreement is effected on the distal level, what need would there be for agreement on the proximal level? Hence what reason is there to assume, say with Quine, that comparison of only proximal rules *ought* to yield determinate translation between idiolects?³⁰

³⁰For further discussion of the relation of proximal assertability rules to truth rules see my "The Price of Correspondence Truth," *Noûs* 20 (1986),

VIII. CAUSES OF VERIFICATIONIST MYOPIA

Given a biological approach, then, there are reasons to think that truth rules may be distal correspondence rules, hence that classical truth conditions may do work for semantics. But Putnam and Dummett claim that any such view is unintelligible. Why?

Although there are passages in both Dummett and Putnam that could be given a less sympathetic reading, the reason is not (or at least is not simply) that these philosophers take understanding to be something that must transpire before consciousness. A more explicit theme is that understanding a language is a practical ability, constituted by a set of *dispositions*, in this case, learned responses: "Now when someone learns a language, what he learns is a practice; he learns to respond, verbally and nonverbally, to utterances and to make utterances of his own" (Dummett);³¹ "language understanding [is] . . . an activity involving 'language entry rules' (procedures for subjecting some sentences to stimulus control), procedures for deductive and inductive inference and 'language exit rules' . . ." (Putnam).³² It follows, Putnam and Dummett now agree, that if a language is characterized by certain truth rules, this fact must be one that shows up in the speech dispositions of the language users. And it follows that if there are no *dispositions* to recognize correspondence truth conditions, sentences can not have correspondence truth conditions.

Putnam's phrase "language entry rules" is a reference to Sellars, but, of course, many other central figures have also held that understanding a language must yield to a dispositional analysis, among them Quine, Davidson, many would say Wittgenstein, and, in the philosophy of mind (re: inner language), the functionalists. Despite this distinguished advocacy, surely Kripke's remark about

pp. 453–468, and also *LTOBC*, especially Chapters 8 and 9. Proximal assertability rules are close relatives of what I there called "intensions." (In this essay I am not emphasizing that perception characteristically is an activity involving overt exploration, a fact that was in the foreground when I spoke of intensions in *LTOBC*. Thus the notion "proximal assertability rules" is a somewhat duller tool than I intended "intensions" to be in *LTOBC*.)

³¹"What is a Theory of Meaning? (II)," p. 82.

³²"Realism and Reason," p. 110.

illegitimate “equation of performance with correctness” is applicable here. To be competent in a language involves that one have a practical grasp of its truth rules. About that everyone agrees. But “true” is clearly a *normative* notion. “True” is how my sincerely uttered sentences are *when they come out right*, when they are, using Kripke’s expression, as they “ought to be,” when I achieve what I purpose in sincerely uttering them. And no mere set of dispositions, no mere performance, determines a measuring “ought,” a standard or norm. No set of dispositions, then, could determine truth rules.

Nor is the normative ingredient in truth provided by the fact that the dispositions that constitute competence in a language must agree with a public norm. Compare games. Consider first a case in which I intend to play the same game as the others do, say, the one they call “chess,” but I mistakenly play by different rules than the others. This is a case of playing wrongly in the sense that I have not played the game I intended, or, perhaps, the one others expected me to. Similarly, if I intend to use the same language as the others, but in fact adopt different truth rules, then I speak wrongly, for I have not spoken the language I intended, or that others expected me to. This is called “not knowing the language” or “making mistakes in the language.” Second, consider a case in which I have no intention to play with the chess pieces as the others do nor do the others expect me to. Then playing by different rules is just playing a different game. It is neither playing chess wrongly nor doing anything else wrongly. The linguistic parallel to this is called “speaking a *different* language.” But speaking wrongly in the sense of speaking *falsely* is still a *third* possibility. Speaking falsely is not just a way of being out of step, nor is it just marching to the beat of a different drummer. Suppose we call it a “rule” of chess that you are supposed to checkmate your opponent. Then speaking falsely is like failing to checkmate the opponent. Better, it is like failing to pick up one straw without moving the others when playing jackstraws. Just as learning the rules better is not the cure for losing at chess or jackstraws, learning the community’s language better is not the cure for bad judgment. And just as whether one succeeds at jackstraws, that is, at not moving the other sticks, does not depend on any agreement with the community, neither does whether one succeeds in speaking

truth in one's language. To purpose to follow certain truth rules is to set a standard for *oneself*—a standard that one may fail to meet.

It is because purposes set standards that “true” is a normative notion and that no set of dispositions could determine truth rules. Similarly, although Dummett and Putnam are right that semantic meaning must be resident somehow in language competence, no set of dispositions equals a competence. First, a disposition does not express a competence unless it is a disposition informed by a purpose. My disposition to fall if left unsupported is no competence, nor is the hoverfly's disposition to chase birds. Conversely, having a competence does not, in general, imply that one has any particular dispositions. If I know how to A—say, to sharpen a drill bit—it doesn't follow that I have a disposition to succeed in A-ing if I try. Perhaps my hands are too cold, or the only grindstone available is not the kind I am practiced at using, or you insist on joggling my elbow. Though I know how to walk, sometimes I trip when I try. Recall the hoverfly, who exhibits a competence whenever he conforms to the distal hoverfly rule in a normal way, yet, due to the inconstancy of conditions outside him, often does not manage to conform to it at all. Nor are there specified conditions under which a person must succeed in order to know how. If I can only sharpen the bit using one sharpening tool whereas you know how to use another, then normal conditions for exercise of my ability to sharpen a drill bit will be different from normal conditions for exercise of yours; each may fail where the other succeeds. Knowing how to do A entails, at best, only that there are *some* normal conditions under which one succeeds in doing A.

Now there is an evident reason why knowing how to A does not, in general, entail having any simple disposition to succeed in A-ing. The reason is that most know-how involves *distal* action, and there is no such thing as a simple *disposition* to involvement with anything distal. How one interacts with things at a distance always depends upon what lies in between, on surrounding conditions. Simple dispositions can concern only reactions to and actions upon that which *touches* one or, perhaps, what is inside one. It follows that to assimilate language competence to a set of dispositions directly begs the question against distal truth rules. There is no need for tortuous arguments to demonstrate that truth rules must then be verificationist. On a dispositional account, to “grasp” cor-

respondence truth rules for each sentence in one's language would be to have a "capacity . . . to evince recognition of the truth of the sentence when and only when the relevant condition is fulfilled" (Dummett).³³ But if a "recognition capacity" is a disposition, it must be a disposition to respond to a proximal stimulus, there being no such thing as a disposition to respond to something distal. And dispositions to respond to proximal stimuli with sentences could correspond, at best, to assertability conditions, certainly not to distal correspondence truth conditions. Q.E.D.

Compare the hoverfly. Assuming that his insides are working right, what he has a *disposition* to do is, at best, to conform to the proximal hoverfly rule. Does it follow that he has no ability to catch females?

It is significant, I think, how close the dispositional view of language understanding is to the more classical view that understanding takes place wholly within consciousness. On the classical view, understanding must ultimately involve relations only to things that touch the mind. On the dispositional view, understanding still involves only what touches the mind or, say, the nervous system. It is easy, then, to slip back and forth between two ways of interpreting the Dummett-Putnam attack upon realist truth.³⁴ Yet what Kripke has shown is that *neither* view of language understanding is a tenable view. Hence, whatever may be said for or against the positive theory of rule following that I have offered, the verificationist vision is surely unnecessarily nearsighted. If Kripke (and Kripke's Wittgenstein) are right, then *whatever* the status of rule following, we have no reason to think that the following of correspondence truth rules is any more *problematic* than is the following of verificationist truth rules.

On the other hand, perhaps what is most puzzling about the following of any kind of language rules is how one could "know" these rules without having a prior language, a prior way of "meaning" or thinking of these rules. Yet surely even the medium of thought, even whatever is currently before the mind or in the

³³"What is a Theory of Meaning? (II)," pp. 80–81.

³⁴I will not attempt to prove that Dummett and Putnam themselves do some sliding, but on Putnam, see *LTOBC*, Epilogue.

head, stands in need of interpretation. Knowing the rules is not a disposition, nor can it be explained in the end by reference to prior representations of the rules. The biological account agrees with both of these considerations.³⁵

University of Connecticut

³⁵Earlier versions of this paper were read at the University of Wisconsin (Madison), Western Michigan University, the University of Maryland, Trinity University, the University of New England (Australia), Australian National University, Monash University and Vanderbilt University. I am grateful to the members of these departments, to Margaret Gilbert and John Troyer, and to unknown referees for *The Philosophical Review*, for helpful comments and suggestions.