

# 10 *A Coherence Theory of Truth and Knowledge*

In this essay I defend what may as well be called a coherence theory of truth and knowledge. The theory I defend is not in competition with a correspondence theory, but depends for its defense on an argument that purports to show that coherence yields correspondence.

The importance of the theme is obvious. If coherence is a test of truth, there is a direct connection with epistemology, for we have reason to believe many of our beliefs cohere with many others, and in that case we have reason to believe many of our beliefs are true. When the beliefs are true, then the primary conditions for knowledge would seem to be satisfied.

Someone might try to defend a coherence theory of truth without defending a coherence theory of knowledge, perhaps on the ground that the holder of a coherent set of beliefs might lack a reason to believe his beliefs coherent. This is not likely, but it may be that someone, though he has true beliefs, and good reasons for holding them, does not appreciate the relevance of reason to belief. Such a one may best be viewed as having knowledge he does not know he has: he thinks he is a skeptic. In a word, he is a philosopher.

Setting aside aberrant cases, what brings truth and knowledge together is meaning. If meanings are given by objective truth conditions, there is a question how we can know that the conditions are satisfied, for this would appear to require a confrontation between what we believe and reality; and the idea of such a confrontation is absurd. But if coherence is a test of truth, then coherence is a test for judging that objective truth conditions are satisfied, and we no longer need to explain meaning on the basis of possible confrontation. My slogan is: correspondence without confrontation. Given a correct epistemology, we can be realists in all departments. We can accept

objective truth conditions as the key to meaning, a realist view of truth, and we can insist that knowledge is of an objective world independent of our thought or language.

Since there is not, as far as I know, a theory that deserves to be called 'the' coherence theory, let me characterize the sort of view I want to defend. It is obvious that not every consistent set of interpreted sentences contains only true sentences, since one such set might contain just the consistent sentence *s* and another just the negation of *s*. And adding more sentences, while maintaining consistency, will not help. We can imagine endless state-descriptions—maximal consistent descriptions—which do not describe our world.

My coherence theory concerns beliefs, or sentences held true by someone who understands them. I do not want to say, at this point, that every possible coherent set of beliefs is true (or contains mostly true beliefs). I shy away from this because it is so unclear what is possible. At one extreme, it might be held that the range of possible maximal sets of beliefs is as wide as the range of possible maximal sets of sentences, and then there would be no point to insisting that a defensible coherence theory concerns beliefs and not propositions or sentences. But there are other ways of conceiving what it is possible to believe which would justify saying not only that all actual coherent belief systems are largely correct but that all possible ones are also. The difference between the two notions of what it is possible to believe depends on what we suppose about the nature of belief, its interpretation, its causes, its holders, and its patterns. Beliefs for me are states of people with intentions, desires, sense organs; they are states that are caused by, and cause, events inside and outside the bodies of their entertainers. But even given all these constraints, there are many things people do believe, and many more that they could. For all such cases, the coherence theory applies.

Of course some beliefs are false. Much of the point of the concept of belief is the potential gap it introduces between what is held to be true and what is true. So mere coherence, no matter how strongly coherence is plausibly defined, cannot guarantee that what is believed is so. All that a coherence theory can maintain is that most of the beliefs in a coherent total set of beliefs are true.

This way of stating the position can at best be taken as a hint, since there is no useful way to count beliefs, and so no clear meaning to the idea that most of a person's beliefs are true. A somewhat better way to put the point is to say there is a presumption in favor

of the truth of a belief that coheres with a significant mass of belief. Every belief in a coherent total set of beliefs is justified in the light of this presumption, much as every intentional action taken by a rational agent (one whose choices, beliefs, and desires cohere in the sense of Bayesian decision theory) is justified. So to repeat, if knowledge is justified true belief, then it would seem that all the true beliefs of a consistent believer constitute knowledge. This conclusion, though too vague and hasty to be right, contains an important core of truth, as I shall argue. Meanwhile I merely note the many problems asking for treatment: What exactly does coherence demand? How much of inductive practice should be included, how much of the true theory (if there is one) of evidential support must be in there? Since no person has a completely consistent body of convictions, coherence with *which* beliefs creates a presumption of truth? Some of these problems will be put in better perspective presently.

It should be clear that I do not hope to define truth in terms of coherence and belief. Truth is beautifully transparent compared to belief and coherence, and I take it as a primitive concept. Truth, as applied to utterances of sentences, shows the disquotational feature enshrined in Tarski's Convention T, and that is enough to fix its domain of application. Relative to a language or a speaker, of course, so there is more to truth than Convention T; there is whatever carries over from language to language or speaker to speaker. What Convention T, and the trite sentences it declares true, like 'Grass is green', spoken by an English speaker, is true if and only if grass is green', reveal is that the truth of an utterance depends on just two things: what the words as spoken mean, and how the world is arranged. There is no further relativism to a conceptual scheme, a way of viewing things, or a perspective. Two interpreters, as unlike in culture, language, and point of view as you please, can disagree over whether an utterance is true, but only if they differ on how things are in the world they share, or what the utterance means.

I think we can draw two conclusions from these simple reflections. First, truth is correspondence with the way things are. (There is no straightforward and nonmisleading way to state this; to get things right, a detour is necessary through the concept of satisfaction in terms of which truth is characterized.<sup>1</sup>) So if a coherence theory of

<sup>1</sup> See my 'True to the Facts', essay 3 in *Inquiries into Truth and Interpretation*, and 'Afterthoughts' in this volume.

truth is acceptable, it must be consistent with a correspondence theory. Second, a theory of knowledge that allows that we can know the truth must be a nonrelativized, noninternal form of realism. So if a coherence theory of knowledge is acceptable, it must be consistent with such a form of realism. My form of realism seems to be neither Hilary Putnam's internal realism nor his metaphysical realism.<sup>2</sup> It is not internal realism because internal realism makes truth relative to a scheme, and this is an idea I do not think is intelligible.<sup>3</sup> A major reason, in fact, for accepting a coherence theory is the unintelligibility of the dualism of a conceptual scheme and a 'world' waiting to be coped with. But my realism is certainly not Putnam's metaphysical realism, for *it* is characterized by being 'radically non-epistemic', which implies that all our best-researched and -established thoughts and theories may be false. I think the independence of belief and truth requires only that *each* of our beliefs may be false. But of course a coherence theory cannot allow that all of them can be wrong.

But why not? Perhaps it is obvious that the coherence of a belief with a substantial body of belief enhances its chance of being true, provided there is reason to suppose the body of belief is true, or largely so. But how can coherence alone supply grounds for belief? Mayhap the best we can do to justify one belief is to appeal to other beliefs. But then the outcome would seem to be that we must accept philosophical skepticism, no matter how unshaken in practice our beliefs remain.

This is skepticism in one of its traditional garbs. It asks: why couldn't all my beliefs hang together and yet be comprehensively false about the actual world? Mere recognition of the fact that it is absurd or worse to try to *confront* our beliefs, one by one, or as a whole, with what they are about does not answer the question nor show the question unintelligible. In short, even a mild coherence theory like mine must provide a skeptic with a reason for supposing coherent beliefs are true. The partisan of a coherence theory can't allow assurance to come from outside the system of belief, while nothing inside can produce support except as it can be shown to rest, finally or at once, on something independently trustworthy.

<sup>2</sup> Hilary Putnam, *Meaning and the Moral Sciences*, 125.

<sup>3</sup> See my 'On the Very Idea of a Conceptual Scheme', essay 13 in *Inquiries into Truth and Interpretation*.

It is natural to distinguish coherence theories from others by reference to the question whether or not justification can or must come to an end. But this does not define the positions, it merely suggests a form the argument may take. For there are coherence theorists who hold that some beliefs can serve as the basis for the rest, while it would be possible to maintain that coherence is not enough, although giving reasons never comes to an end. What distinguishes a coherence theory is simply the claim that nothing can count as a reason for holding a belief except another belief. Its partisan rejects as unintelligible the request for a ground or source of justification of another ilk. As Rorty has put it, 'nothing counts as justification unless by reference to what we already accept, and there is no way to get outside our beliefs and our language so as to find some test other than coherence'.<sup>4</sup> About this I am, as you see, in agreement with Rorty. Where we differ, if we do, is on whether there remains a question how, given that we cannot 'get outside our beliefs and our language so as to find some test other than coherence', we nevertheless can have knowledge of, and talk about, an objective public world which is not of our own making. I think this question does remain, while I suspect that Rorty doesn't think so. If this is his view, then he must think I am making a mistake in trying to answer the question. Nevertheless, here goes.

It will promote matters at this point to review very hastily some of the reasons for abandoning the search for a basis for knowledge outside the scope of our beliefs. By 'basis' here I mean specifically an epistemological basis, a source of justification.

The attempts worth taking seriously attempt to ground belief in one way or another on the testimony of the senses: sensation, perception, the given, experience, sense data, the passing show. All such theories must explain at least these two things: what, exactly, is the relation between sensation and belief that allows the first to justify the second? and, why should we believe our sensations are reliable, that is, why should we trust our senses?

The simplest idea is to identify certain beliefs with sensations. Thus Hume seems not to have distinguished between perceiving a green spot and perceiving that a spot is green. (An ambiguity in the word 'idea' was a great help here.) Other philosophers noted Hume's confusion, but tried to attain the same results by reducing the gap

<sup>4</sup> Richard Rorty, *Philosophy and the Mirror of Nature*, 178.

between perception and judgement to zero by attempting to formulate judgements that do not go beyond stating that the perception or sensation or presentation exists (whatever that may mean). Such theories do not justify beliefs on the basis of sensations, but try to justify certain beliefs by claiming that they have exactly the same epistemic content as a sensation. There are two difficulties with such a view: first, if the basic beliefs do not exceed in content the corresponding sensation, they cannot support any inference to an objective world; and second, there are no such beliefs.

A more plausible line is to claim that we cannot be wrong about how things appear to us to be. If we believe we have a sensation, we do; this is held to be an analytic truth, or a fact about how language is used.

It is difficult to explain this supposed connection between sensations and some beliefs in a way that does not invite skepticism about other minds, and in the absence of an adequate explanation, there should be a doubt about the implications of the connection for justification. But in any case, it is unclear how, on this line, sensations justify the belief in those sensations. The point is rather that such beliefs require no justification, for the existence of the belief entails the existence of the sensation, and so the existence of the belief entails its own truth. Unless something further is added, we are back to another form of coherence theory.

Emphasis on sensation or perception in matters epistemological springs from the obvious thought: sensations are what connect the world and our beliefs, and they are candidates for justifiers because we often are aware of them. The trouble we have been running into is that the justification seems to depend on the awareness, which is just another belief.

Let us try a bolder tack. Suppose we say that sensations themselves, verbalized or not, justify certain beliefs that go beyond what is given in sensation. So, under certain conditions, having the sensation of seeing a green light flashing may justify the belief that a green light is flashing. The problem is to see how the sensation justifies the belief. Of course if someone has the sensation of seeing a green light flashing, it is likely, under certain circumstances, that a green light is flashing. We can say this, since we know of his sensation, but *he* can't say it, since we are supposing he is justified without having to depend on believing he has the sensation. Suppose he believed he didn't have the sensation. Would the sensation still justify him in the belief in an objective flashing green light?

The relation between a sensation and a belief cannot be logical, since sensations are not beliefs or other propositional attitudes. What then is the relation? The answer is, I think, obvious: the relation is causal. Sensations cause some beliefs and in *this* sense are the basis or ground of those beliefs. But a causal explanation of a belief does not show how or why the belief is justified.

The difficulty of transmuting a cause into a reason plagues the antioherentist again if he tries to answer our second question: what justifies the belief that our senses do not systematically deceive us? For even if sensations justify belief in sensation, we do not yet see how they justify belief in external events and objects.

According to Quine, science tells us that 'our only source of information about the external world is through the impact of light rays and molecules upon our sensory surfaces'.<sup>5</sup> What worries me is how to read the words 'source' and 'information'. Certainly it is true that events and objects in the external world cause us to believe things about the external world, and much, if not all, of the causality takes a route through the sense organs. The notion of information, however, applies in a nonmetaphorical way only to the engendered beliefs. So 'source' has to be read simply as 'cause' and 'information' as 'true belief' or 'knowledge'. Justification of beliefs caused by our senses is not yet in sight.

The approach to the problem of justification we have been tracing must be wrong. We have been trying to see it this way: a person has all his beliefs about the world—that is, all his beliefs. How can he tell if they are true, or apt to be true? This is possible, we have been assuming, only by connecting his beliefs to the world, confronting

<sup>5</sup> W. V. Quine, 'The Nature of Natural Knowledge', 68. Many other passages in Quine suggest that he hopes to assimilate sensory causes to evidence. In *Word and Object*, 22, he writes that 'surface irritations . . . exhaust our clues to an external world'. In *Ontological Relativity*, 75, we find that 'The stimulation of his sensory receptors is all the evidence anybody has had to go on, ultimately, in arriving at his picture of the world.' On the same page: 'Two cardinal tenets of empiricism remain unassailable . . . One is that whatever evidence there is for science is sensory evidence. The other . . . is that all inculcation of meanings of words, must rest ultimately on sensory evidence.' In *The Roots of Reference*, 37–8, Quine says 'observations' are basic 'both in the support of theory and in the learning of language', and then goes on, 'What are observations? They are visual, auditory, tactual, olfactory. They are sensory, evidently, and thus subjective . . . Should we say then that the observation is not the sensation . . . ? No . . .'. Quine goes on to abandon talk of observations in favor of talk of observation sentences. But of course observation sentences, unlike observations, cannot play the role of evidence unless we have reason to believe they are true.

certain of his beliefs with the deliverances of the senses one by one, or perhaps confronting the totality of his beliefs with the tribunal of experience. No such confrontation makes sense, for of course we can't get outside our skins to find out what is causing the internal happening of which we are aware. Introducing intermediate steps or entities into the causal chain, like sensations or observations, serves only to make the epistemological problem more obvious. For if the intermediaries are merely causes, they don't justify the beliefs they cause, while if they deliver information, they may be lying. The moral is obvious. Since we can't swear intermediaries to truthfulness, we should allow no intermediaries between our beliefs and their objects in the world. Of course there are causal intermediaries. What we must guard against are epistemic intermediaries.

There are common views of language that encourage bad epistemology. This is no accident, of course, since theories of meaning are connected with epistemology through attempts to answer the question how one determines that a sentence is true. If knowing the meaning of a sentence (knowing how to give a correct interpretation of it) involves, or is, knowing how it could be recognized to be true, then the theory of meaning raises the same question we have been struggling with, for giving the meaning of a sentence will demand that we specify what would justify asserting it. Here the coherentist will hold that there is no use looking for a source of justification outside of other sentences held true, while the foundationalist will seek to anchor at least some words or sentences to non-verbal rocks. This view is held, I think, both by Quine and by Michael Dummett.

Dummett and Quine differ, to be sure. In particular, they disagree about holism, the claim that the truth of our sentences must be tested together rather than one by one. And they disagree also, and consequently, about whether there is a useful distinction between analytic and synthetic sentences, and about whether a satisfactory theory of meaning can allow the sort of indeterminacy Quine argues for. (On all these points, I am Quine's faithful student.)

But what concerns me here is that Quine and Dummett agree on a basic principle, which is that whatever there is to meaning must be traced back somehow to experience, the given, or patterns of sensory stimulation, something intermediate between belief and the usual objects our beliefs are about. Once we take this step, we open the door to skepticism, for we must then allow that a very great many—perhaps most—of the sentences we hold to be true may in fact be

false. It is ironical. Trying to make meaning accessible has made truth inaccessible. When meaning goes epistemological in this way, truth and meaning are necessarily divorced. One can, of course, arrange a shotgun wedding by redefining truth as what we are justified in asserting. But this does not marry the original mates.

Take Quine's proposal that whatever there is to the meaning (information value) of an observation sentence is determined by the patterns of sensory stimulation that would cause a speaker to assent to or dissent from the sentence. This is a marvelously ingenious way of capturing what is appealing about verificationist theories without having to talk of meanings, sense data, or sensations; for the first time it made plausible the idea that one could, and should, do what I call the theory of meaning without need of what Quine calls meanings. But Quine's proposal, like other forms of verificationism, makes for skepticism. For clearly a person's sensory stimulations could be just as they are and yet the world outside very different. (Remember the brain in the vat.)

Quine's way of doing without meanings is subtle and complicated. He ties the meanings of some sentences directly to patterns of stimulation (which also constitute the evidence, Quine thinks, for assenting to the sentence), but the meanings of further sentences are determined by how they are conditioned to the original, or observation, sentences. The facts of such conditioning do not permit a sharp division between sentences held true by virtue of meaning and sentences held true on the basis of observation. Quine made this point by showing that if one way of interpreting a speaker's utterances was satisfactory, so were many others. This doctrine of the indeterminacy of translation, as Quine called it, should be viewed as neither mysterious nor threatening. It is no more mysterious than the fact that temperature can be measured in centigrade or Fahrenheit (or any linear transformation of those numbers). And it is not threatening because the very procedure that demonstrates the degree of indeterminacy at the same time demonstrates that what is determinate is all we need.

In my view, erasing the line between the analytic and synthetic saved philosophy of language as a serious subject by showing how it could be pursued without what there cannot be: determinate meanings. I now suggest also giving up the distinction between observation sentences and the rest. For the distinction between sentences belief in whose truth is justified by sensations and sentences belief in

whose truth is justified only by appeal to other sentences held true is as anathema to the coherentist as the distinction between beliefs justified by sensations and beliefs justified only by appeal to further beliefs. Accordingly, I suggest we give up the idea that meaning or knowledge is grounded on something that counts as an ultimate source of evidence. No doubt meaning and knowledge depend on experience, and experience ultimately on sensation. But this is the 'depend' of causality, not of evidence or justification.

I have now stated my problem as well as I can. The search for an empirical foundation for meaning or knowledge leads to skepticism, while a coherence theory seems at a loss to provide any reason for a believer to believe that his beliefs, if coherent, are true. We are caught between a false answer to the skeptic, and no answer.

The dilemma is not a true one. What is needed to answer the skeptic is to show that someone with a (more or less) coherent set of beliefs has a reason to suppose his beliefs are not mistaken in the main. What we have shown is that it is absurd to look for a justifying ground for the totality of beliefs, something outside this totality which we can use to test or compare with our beliefs. The answer to our problem must then be to find a *reason* for supposing most of our beliefs are true that is not a form of *evidence*.

My argument has two parts. First I urge that a correct understanding of the speech, beliefs, desires, intentions, and other propositional attitudes of a person leads to the conclusion that most of a person's beliefs must be true, and so there is a legitimate presumption that any one of them, if it coheres with most of the rest, is true. Then I go on to claim that anyone with thoughts, and so in particular anyone who wonders whether he has any reason to suppose he is generally right about the nature of his environment, must know what a belief is, and how in general beliefs are to be detected and interpreted. These being perfectly general facts we cannot fail to use when we communicate with others, or when we try to communicate with others, or even when we merely think we are communicating with others, there is a pretty strong sense in which we can be said to know that there is a presumption in favor of the overall truthfulness of anyone's beliefs, including our own. So it is bootless for someone to ask for some *further* reassurance; that can only add to his stock of beliefs. All that is needed is that he recognize that belief is in its nature veridical.

Belief can be seen to be veridical by considering what determines the existence and contents of a belief. Belief, like the other so-called

propositional attitudes, is supervenient on facts of various sorts, behavioral, neurophysiological, biological, and physical. The reason for pointing this out is not to encourage definitional or nomological reduction of psychological phenomena to something more basic, and certainly not to suggest epistemological priorities. The point is rather understanding. We gain one kind of insight into the nature of the propositional attitudes when we relate them systematically to one another and to phenomena on other levels. Since the propositional attitudes are deeply interlocked, we cannot learn the nature of one by first winning understanding of another. As interpreters, we work our way into the whole system, depending much on the pattern of inter-relationships.

Take, for example, the interdependence of belief and meaning. What a sentence means depends partly on the external circumstances that cause it to win some degree of conviction; and partly on the relations, grammatical or logical, that the sentence has to other sentences held true with varying degrees of conviction. Since these relations are themselves translated directly into beliefs, it is easy to see how meaning depends on belief. Belief, however, depends equally on meaning, for the only access to the fine structure and individuation of beliefs is through the sentences speakers and interpreters of speakers use to express and describe beliefs. If we want to illuminate the nature of meaning and belief, therefore, we need to start with something that assumes neither. Quine's suggestion, which I shall essentially follow, is to take *prompted assent* as basic, the causal relation between assenting to a sentence and the cause of such assent. This is a fair place to start the project of identifying beliefs and meanings, since a speaker's assent to a sentence depends both on what he means by the sentence and on what he believes about the world. Yet it is possible to know that a speaker assents to a sentence without knowing either what the sentence, as spoken by him, means, or what belief is expressed by it. Equally obvious is the fact that once an interpretation has been given for a sentence assented to, a belief has been attributed. If correct theories of interpretation are not unique (do not lead to uniquely correct interpretations), the same will go for attributions of belief, of course, as tied to acquiescence in particular sentences.

A speaker who wishes his words to be understood cannot systematically deceive his would-be interpreters about when he assents to sentences—that is, holds them true. As a matter of principle, then,

meaning, and by its connection with meaning, belief also, are open to public determination. I shall take advantage of this fact in what follows and adopt the stance of a radical interpreter when asking about the nature of belief. What a fully informed interpreter could learn about what a speaker means is all there is to learn; the same goes for what the speaker believes.<sup>6</sup>

The interpreter's problem is that what he is assumed to know—the causes of assents to sentences of a speaker—is, as we have seen, the product of two things he is assumed not to know, meaning and belief. If he knew the meanings he would know the beliefs, and if he knew the beliefs expressed by sentences assented to, he would know the meanings. But how can he learn both at once, since each depends on the other?

The general lines of the solution, like the problem itself, are owed to Quine. I will, however, introduce some changes into Quine's solution, as I have into the statement of the problem. The changes are directly relevant to the issue of epistemological skepticism.

Radical interpretation (which is much, but not entirely, like Quine's radical translation) aims at producing a Tarski-style characterization of truth for the speaker's language, and a theory of his beliefs. (The second follows from the first plus the presupposed knowledge of sentences held true.) This adds little to Quine's program of translation, since translation of the speaker's language into one's own plus a theory of truth for one's own language add up to a theory of truth for the speaker. But the shift to the semantic notion of truth from the syntactic notion of translation puts the formal restrictions of a theory of truth in the foreground, and emphasizes one aspect of the close relation between truth and meaning.

The principle of charity plays a crucial role in Quine's method, and an even more crucial role in my variant. In either case, the principle directs the interpreter to translate or interpret so as to read some of his own standards of truth into the pattern of sentences held true by the speaker. The point of the principle is to make the speaker intelligible, since too great deviations from consistency and correctness leave no common ground on which to judge either conformity or difference. From a formal point of view, the principle of charity helps

<sup>6</sup> I now think it is essential, in doing radical interpretation, to include the desires of the speaker from the start, so that the springs of action and intention, both belief and desire, are related to meaning. But in the present essay it is not necessary to introduce this further factor.

solve the problem of the interaction of meaning and belief by restraining the degrees of freedom allowed belief while determining how to interpret words.

We have no choice, Quine has urged, but to read our own logic into the thoughts of a speaker; Quine says this for the sentential calculus, and I would add the same for first-order quantification theory. This leads directly to the identification of the logical constants, as well as to the assignment of a logical form to each sentence.

Something like charity operates in the interpretation of those sentences whose causes of assent come and go with time and place: when the interpreter finds a sentence of the speaker the speaker assents to regularly under conditions the interpreter recognizes, the interpreter takes those conditions to be the truth conditions of the speaker's sentence. This is only roughly right, as we shall see in a moment.

Sentences and predicates less directly geared to easily detected goings-on can, in Quine's canon, be interpreted at will, given only the constraints of interconnections, with sentences conditioned directly to the world. Here I would extend the principle of charity to favor interpretations that as far as possible preserve truth: I think it makes for mutual understanding, and hence for better interpretation, to interpret what the speaker accepts as true as true when we can. In this matter, I have less choice than Quine, because I do not see how to draw the line between observation sentences and theoretical sentences at the start. There are several reasons for this, but the one most relevant to the present topic is that this distinction is ultimately based on an epistemological consideration of a sort I have renounced: observation sentences are directly based on something like sensation—patterns of sensory stimulation—and this is an idea I have been urging leads to skepticism. Without the direct tie to sensation or stimulation, the distinction between observation sentences and others can't be drawn on epistemologically significant grounds. The distinction between sentences whose causes to assent come and go with observable circumstances and those a speaker clings to through change remains, however, and offers the possibility of interpreting the words and sentences beyond the logical.

The details are not here to the point. What should be clear is that if the account I have given of how belief and meaning are related and understood by an interpreter is right, then most of the sentences a



speaker holds to be true—especially the ones he holds to most stubbornly, the ones most central to the system of his beliefs—most of these sentences *are* true, at least in the opinion of the interpreter. For the only, and therefore unimpeachable, method available to the interpreter automatically puts the speaker's beliefs in accord with the standards of logic of the interpreter, and hence credits the speaker with the plain truths of logic. Needless to say, there are degrees of logical and other consistency, and perfect consistency is not to be expected. What needs emphasis is only the methodological necessity for finding consistency enough.

Analogously, it is impossible for an interpreter to understand a speaker and at the same time discover the speaker to be largely wrong about the world. For the interpreter interprets sentences held true (which is not to be distinguished from attributing beliefs) according to the events and objects in the outside world that cause the sentence to be held true.

What I take to be the important aspect of this approach is apt to be missed because the approach reverses our natural way of thinking of communication derived from situations in which understanding has already been secured. Once understanding has been secured, we are able, often, to learn what a person believes quite independently of what caused him to believe it. This may lead us to the crucial, indeed fatal, conclusion that we can in general fix what someone means independently of what he believes and independently of what caused the belief. But if I am right, we can't in general first identify beliefs and meanings and then ask what caused them. The causality plays an indispensable role in determining the content of what we say and believe. This is a fact we can be led to recognize by taking up, as we have, the interpreter's point of view.

It is an artifact of the interpreter's correct interpretation of a person's speech and attitudes that there is a large degree of truth and consistency in the thought and speech of an agent. But this is truth and consistency by the interpreter's standards. Why couldn't it happen that speaker and interpreter understand one another on the basis of shared but erroneous beliefs? This can, and no doubt often does, happen. But it cannot be the rule. For imagine for a moment an interpreter who is omniscient about the world, and about what does and would cause a speaker to assent to any sentence in his (potentially unlimited) repertoire. The omniscient interpreter, using the same method as the fallible interpreter, finds the fallible speaker

largely consistent and correct. By his own standards, of course, but since these are objectively correct, the fallible speaker is seen to be largely correct and consistent by objective standards. We may also, if we want, let the omniscient interpreter turn his attention to the fallible interpreter of the fallible speaker. It turns out that the fallible interpreter can be wrong about some things, but not in general; and so he cannot share universal error with the agent he is interpreting. Once we agree to the general method of interpretation I have sketched, it becomes impossible correctly to hold that anyone could be mostly wrong about how things are.

There is, as I noted above, a key difference between the method of radical interpretation I am now recommending, and Quine's method of radical translation. The difference lies in the nature of the choice of causes that govern interpretation. Quine makes interpretation depend on patterns of sensory stimulation, while I make it depend on the external events and objects the sentence is interpreted as being about. Thus Quine's notion of meaning is tied to sensory criteria, something he thinks can be treated also as evidence. This leads Quine to give epistemic significance to the distinction between observation sentences and others, since observation sentences are supposed, by their direct conditioning to the senses, to have a kind of extralinguistic justification. This is the view against which I argued in the first part of my essay, urging that sensory stimulations are indeed part of the causal chain that leads to belief, but cannot, without confusion, be considered to be evidence, or a source of justification, for the stimulated beliefs.

What stands in the way of global skepticism of the senses is, in my view, the fact that we must, in the plainest and methodologically most basic cases, take the objects of a belief to be the causes of that belief. And what we, as interpreters, must take them to be is what they in fact are. Communication begins where causes converge: your utterance means what mine does if belief in its truth is systematically caused by the same events and objects.<sup>7</sup>

The difficulties in the way of this view are obvious, but I think they can be overcome. The method applies directly, at best, only to

<sup>7</sup> It is clear that the causal theory of meaning has little in common with the causal theories of reference of Kripke and Putnam. Those theories look to causal relations between names and objects of which speakers may well be ignorant. The chance of systematic error is thus increased. My causal theory does the reverse by connecting the cause of a belief with its object.



occasion sentences—the sentences assent to which is caused systematically by common changes in the world. Further sentences are interpreted by their conditioning to occasion sentences, and the appearance in them of words that appear also in occasion sentences. Among occasion sentences, some will vary in the credence they command not only in the face of environmental change, but also in the face of change of credence awarded related sentences. Criteria can be developed on this basis to distinguish degrees of observability on internal grounds, without appeal to the concept of a basis for belief outside the circle of beliefs.

Related to these problems, and easier still to grasp, is the problem of error. For even in the simplest cases it is clear that the same cause (a rabbit scampers by) may engender different beliefs in speaker and observer, and so encourage assent to sentences which cannot bear the same interpretation. It is no doubt this fact that made Quine turn from rabbits to patterns of stimulation as the key to interpretation. Just as a matter of statistics, I'm not sure how much better one approach is than the other. Is the relative frequency with which identical patterns of stimulation will touch off assent to 'Gavagai' and 'Rabbit' greater than the relative frequency with which a rabbit touches off the same two responses in speaker and interpreter? Not an easy question to test in a convincing way. But let the imagined results speak for Quine's method. Then I must say, what I must say in any case, the problem of error cannot be met sentence by sentence, even at the simplest level. The best we can do is cope with error holistically, that is, we interpret so as to make an agent as intelligible as possible, given his actions, his utterances, and his place in the world. About some things we will find him wrong, as the necessary cost of finding him elsewhere right. As a rough approximation, finding him right means identifying the causes with the objects of his beliefs, giving special weight to the simplest cases, and countenancing error where it can be best explained.

Suppose I am right that an interpreter must so interpret as to make a speaker or agent largely correct about the world. How does this help the person himself who wonders what reason he has to think his beliefs are mostly true? How can he learn about the causal relations between the real world and his beliefs that lead the interpreter to interpret him as being on the right track?

The answer is contained in the question. In order to doubt or wonder about the provenance of his beliefs, an agent must know

what belief is. This brings with it the concept of objective truth, for the notion of a belief is the notion of a state that may or may not jibe with reality. But beliefs are also identified, directly and indirectly, by their causes. What an omniscient interpreter knows a fallible interpreter gets right enough if he understands a speaker, and this is just the complicated causal truth that makes us the believers we are, and fixes the contents of our beliefs. The agent has only to reflect on what a belief is to appreciate that most of his basic beliefs are true, and among his beliefs, those most securely held and that cohere with the main body of his beliefs are the most apt to be true. The question 'how do I know my beliefs are generally true?' thus answers itself, simply because beliefs are by nature generally true. Rephrased or expanded, the question becomes, 'How can I tell whether my beliefs, which are by their nature generally true, are generally true?'

All beliefs are justified in this sense: they are supported by numerous other beliefs (otherwise they wouldn't be the beliefs they are), and have a presumption in favor of their truth. The presumption increases the larger and more significant the body of beliefs with which a belief coheres and, there being no such thing as an isolated belief, there is no belief without a presumption in its favor. In this respect, interpreter and interpreted differ. From the interpreter's point of view, methodology enforces a general presumption of truth for the body of beliefs as a whole, but the interpreter does not need to presume each particular belief of someone else is true. The general presumption applied to others does not make them globally right, as I have emphasized, but provides the background against which to accuse them of error. But from each person's own vantage point, there must be a graded presumption in favor of each of his own beliefs.

We cannot, alas, draw the picturesque and pleasant conclusion that all true beliefs constitute knowledge. For though all of a believer's beliefs are to some extent justified to him, some may not be justified enough, or in the right way, to constitute knowledge. The general presumption in favor of the truth of belief serves to rescue us from a standard form of skepticism by showing why it is impossible for all our beliefs to be false together. This leaves almost untouched the task of specifying the conditions of knowledge. I have not been concerned with the canons of evidential support (if such there be), but to show that all that counts as evidence or justification for a belief must come from the same totality of belief to which it belongs.

## Afterthoughts

A few aging philosophes, which category may include Quine, Putnam, and Dummett, and certainly includes me, are still puzzling over the nature of truth and its connections or lack of connections with meaning and epistemology. Rorty thinks we should stop worrying; he believes philosophy has seen through or outgrown the puzzles and should turn to less heavy and more interesting matters. He is particularly impatient with me for not conceding that the old game is up because he finds in my work useful support for his enlightened stance; underneath my 'out-dated rhetoric' he detects the outlines of a largely correct attitude.

In 'Pragmatism, Davidson, and Truth' Rorty urges two things: that my view of truth amounts to a rejection of both coherence and correspondence theories and should properly be classed as belonging to the pragmatist tradition, and that I should not pretend that I am answering the skeptic when I am really telling him to get lost. I pretty much concur with him on both points.

In our 1983 discussion at the Pacific Division Meeting of the American Philosophical Association I agreed to stop calling my position either a coherence or a correspondence theory if he would give up the pragmatist theory of truth. He has done his part; he now explicitly rejects both James and Peirce on truth. I am glad to hold to my side of the bargain. If it had not already been published, I would now change the title of 'A Coherence Theory of Truth and Knowledge', and I would not describe the project as showing how 'coherence yields correspondence'. On internal evidence alone, as Rorty points out, my view cannot be called a correspondence theory. As long ago as 1969 ('True to the Facts'<sup>8</sup>) I argued that nothing can usefully and intelligibly be said to correspond to a sentence; and I repeated this in 'A Coherence Theory of Truth and Knowledge'. I thought then the fact that in characterizing truth for a language it is necessary to put words into relation with objects was enough to give some grip for the idea of correspondence; but this now seems to me a mistake. The mistake is in a way only a misnomer, but terminological infelicities have a way of breeding conceptual confusion, and so it is here. Correspondence theories have always been conceived as

providing an explanation or analysis of truth, and this a Tarski-style theory of truth certainly does not do. I would also now reject the point generally made against correspondence theories that there is no way we could ever tell whether our sentences or beliefs correspond to reality. This criticism is at best misleading, since no one has ever explained in what such a correspondence could consist; and, worse, it is predicated on the false assumption that truth is transparently epistemic.

I also regret having called my view a 'coherence theory'. My emphasis on coherence was properly just a way of making a negative point, that 'all that counts as evidence or justification for a belief must come from the same totality of belief to which it belongs'. Of course this negative claim has typically led those philosophers who held it to conclude that reality and truth are constructs of thought; but it does not lead me to this conclusion, and for this reason if no other I ought not to have called my view a coherence theory. There is also a less weighty reason for not stressing coherence. Coherence is nothing but consistency. It is certainly in favor of a set of beliefs that they are consistent, but there is no chance that a person's beliefs will not tend to be consistent, since beliefs are individuated in part by their logical properties; what is not largely consistent with many other beliefs cannot be identified as a belief. The main thrust of 'A Coherence Theory of Truth and Knowledge' has little to do with consistency; the important thesis for which I argue is that belief is intrinsically veridical. This is the ground on which I maintain that while truth is not an epistemic concept, neither is it wholly severed from belief (as it is in different ways by both correspondence and coherence theories).

My emphasis on coherence was misplaced; calling my view a 'theory' was a plain blunder. In his paper Rorty stressed a minimalist attitude towards truth that he correctly thought we shared. It could be put this way: truth is as clear and basic a concept as we have. Tarski has given us an idea of how to apply the general concept (or try to apply it) to particular languages on the assumption that we already understand it; but of course he didn't show how to define it in general (he proved, rather, that this couldn't be done). Any further attempt to explain, define, analyze, or explicate the concept will be empty or wrong: correspondence theories, coherence theories, pragmatist theories, theories that identify truth with warranted assertability (perhaps under 'ideal' or 'optimum' conditions), theories that ask

<sup>8</sup> Repr. in *Inquiries into Truth and Interpretation*.

truth to explain the success of science, or serve as the ultimate outcome of science or the conversations of some élite, all such theories either add nothing to our understanding of truth or have obvious counterexamples. Why on earth should we expect to be able to reduce truth to something clearer or more fundamental? After all, the only concept Plato succeeded in defining was mud (dirt and water). Putnam's comparison of various attempts to characterize truth with the attempts to define 'good' in naturalistic terms seems to me, as it does to Rorty, apt. It also seems to apply to Putnam's identification of truth with idealized warranted assertability.<sup>9</sup>

A theory of truth for a speaker, or group of speakers, while not a definition of the general concept of truth, does give a firm sense of what the concept is good for; it allows us to say, in a compact and clear way, what someone who understands that speaker, or those speakers, knows. Such a theory also invites the question how an interpreter could confirm its truth—a question which without the theory could not be articulated. The answer will, as I try to show in 'A Coherence Theory of Truth and Knowledge', bring out essential relations among the concepts of meaning, truth, and belief. If I am right, each of these concepts requires the others, but none is subordinate to, much less definable in terms of, the others. Truth emerges not as wholly detached from belief (as a correspondence theory would make it) nor as dependent on human methods and powers of discovery (as epistemic theories of truth would make it). What saves truth from being 'radically non-epistemic' (in Putnam's words) is not that truth is epistemic but that belief, through its ties with meaning, is intrinsically veridical.

Finally, how about Rorty's admonition to stop trying to answer the skeptic, and tell him to get lost? A short response would be that the skeptic has been told this again and again over the millennia and never seems to listen; like the philosopher he is, he wants an argument. To spell this out a bit: there is perhaps the suggestion in Rorty's 'Pragmatism, Davidson, and Truth' that a 'naturalistic' approach to the problems of meaning and the propositional attitudes will automatically leave the skeptic no room for maneuver. This thought, whether or not it is Rorty's, is wrong. Quine's naturalized epistemology, because it is based on the empiricist premise that what we mean and what we think is conceptually (and not merely causally) founded

on the testimony of the senses, is open to standard skeptical attack. I was much concerned in 'A Coherence Theory of Truth and Knowledge' to argue for an alternative approach to meaning and knowledge, and to show that if this alternative were right, skepticism could not get off the ground. I agree with Rorty to this extent; I set out not to 'refute' the skeptic, but to give a sketch of what I think to be a correct account of the foundations of linguistic communication and its implications for truth, belief, and knowledge. If one grants the correctness of this account, one can tell the skeptic to get lost.

Where Rorty and I differ, if we do, is in the importance we attach to the arguments that lead to the skeptic's undoing, and in the interest we find in the consequences for knowledge, belief, truth, and meaning. Rorty wants to dwell on where the arguments have led: to a position which allows us to dismiss the skeptic's doubts, and so to abandon the attempt to provide a general justification for knowledge claims—a justification that is neither possible nor needed. Rorty sees the history of Western philosophy as a confused and victorious battle between unintelligible skepticism and lame attempts to answer it. Epistemology from Descartes to Quine seems to me just one complex, and by no means unilluminating, chapter in the philosophical enterprise. If that chapter is coming to a close, it will be through recourse to modes of analysis and adherence to standards of clarity that have always distinguished the best philosophy, and will, with luck and enterprise, continue to do so.

<sup>9</sup> Hilary Putnam, *Realism and Reason*, p. xviii.