

W. V. Quine (1908–2000)

PETER HYLTON

Willard Van Orman Quine was born on June 25, 1908. He was graduated from Oberlin College with a degree in mathematics, *summa cum laude*, in 1930; his senior honors thesis was a proof within the system of Whitehead and Russell's *Principia Mathematica* (1910–13), which he studied largely without aid from his teachers. Whitehead was in the Philosophy Department at Harvard, so it was there that Quine went to do graduate work, although Whitehead was no longer teaching logic and Quine had done little undergraduate work in philosophy. Nonetheless, he completed a Ph.D. in two years, graduating in 1932. His dissertation generalized *Principia's* treatment of classes so that it included dyadic relations, instead of treating the latter separately. Along the way, Quine clarified and reformulated the basis of the system – a point to which we shall return.

Quine spent the academic year 1932–3 in Europe on a Sheldon Fellowship. He spent five months in Vienna, attending some meetings of the Vienna Circle. More important, perhaps, was a shorter stay in Prague, where he had extensive conversations with Rudolf Carnap, then completing *The Logical Syntax of Language* (Carnap 1934). While in Europe he was elected as one of the first group of Junior Fellows in Harvard's newly formed Society of Fellows: a three-year fellowship, without teaching obligations. He spent most of these years working on logic and set theory, though some of it on other aspects of philosophy. Under this latter head, he gave three lectures on Carnap, essentially expounding, in a strongly approving fashion, what he took to be the doctrines of *The Logical Syntax of Language*. (The text of these lectures is now published in Creath 1990.) In 1936 Quine became Faculty Instructor at Harvard. Except for service in the US Navy during World War II, he held faculty positions there from that time until his retirement in 1978. He remained philosophically active and engaged for twenty years after retirement, continuing to write and publish into his nineties.

It will be helpful to put Quine's work in the context of what I shall call twentieth-century scientific philosophy, a movement within the broader stream of twentieth-century analytic philosophy. Key figures in twentieth-century scientific philosophy (other than Quine) include Bertrand Russell and Rudolf Carnap, as well as others often identified as logical positivists or logical empiricists; Frege and Wittgenstein also made crucial contributions to the movement.

Let us try briefly to characterize this movement by aims and doctrines rather than by its participants. Perhaps most notable is the emphasis on knowledge, and its objects, rather than on ethics or politics or aesthetics or history or the human condition, as the primary concern of philosophy; an emphasis, one might say, on the True rather than on the Good or the Beautiful. This emphasis is equally an emphasis on science, especially on the natural sciences. It is characteristic of scientific philosophy to take the natural sciences as paradigmatic of all knowledge. Part of this view is the doctrine that the Vienna Circle called “the Unity of Science.” The point here is the unity of all real knowledge, for the German word *Wissenschaft* is broader in its scope than most current uses of its invariable English translation, “science.” (Quine, however, makes it explicit that he uses the word “science” broadly; see Quine 1995: 49.) According to this view, there are no fundamental divisions of aim or method among the various branches of knowledge. Along with this, there is a suspicion, or worse, of the claims of metaphysics, and of any claims neither answerable in straightforward fashion to the findings of empirical science nor provable by logic and mathematics.

This brief sketch at once raises questions about there being any role for philosophy, even that of the scientific philosophers themselves. A central idea here was that philosophy is not to add to our knowledge but is, rather, to *analyze* the knowledge that the sciences give us, and thereby to give us greater clarity about that knowledge and its basis. The tool of this analysis was, above all, logic: the logic of Frege and Russell (see FREGE and RUSSELL). This logic held out an ideal of clarity; one aspect of the philosopher’s task was to impose a similar clarity upon other subjects. This line of thought suggests an assimilation of philosophy to logic, but does not by itself account for the possibility of either of these subjects. Russell sought to do this by postulating an a priori insight, which might strike some as a large concession to metaphysics, in the pejorative sense. Carnap, drawing on the early work of Wittgenstein, held that logic is analytic, empty of content, and hence not genuine knowledge at all. Philosophy too makes no claims about the world. Its analyses of language simply make explicit what is already there; it recommends a certain kind of language for this or that scientific purpose, but a recommendation is not a claim, and is presumably not in need of the same sort of justification. This emphasis on language is connected with Carnap’s view that analytic truths are true in virtue of language, true by virtue of the meanings of the words making them up. Given the importance of analyticity as accounting for logic, for mathematics, and for philosophy itself, this throws an enormous explanatory burden on the notion of language.

We have just very briefly sketched the tradition of twentieth-century scientific philosophy. Quine’s position relative to this tradition is ambivalent. On the one hand, he is its greatest exponent in the last forty years of the century. On the other hand he revolutionizes it, in such a way that one might say that he rejects the tradition rather than continuing it. Both Russell and Carnap attributed great importance to the natural sciences but nevertheless held that logic, mathematics, and philosophy itself, all have a status that is quite different from that of, say, physics or chemistry, or history or sociology. The former are independent of observation, and thus a priori – however exactly that idea is to be understood – while the latter are a posteriori, empirical and ultimately answerable to observation and sensory experience. Quine rejects the

idea that there is a fundamental epistemological distinction here. This rejection – which Quine himself sometimes speaks of as his *naturalism* – is fundamental for his philosophy. We need to see why he rejects the a priori, and how he can get by without it; we shall then begin to show how his general approach to philosophy flows, in large part, from this step.

Analyticity and the a priori

Like Carnap, Quine rejects any idea of the a priori as based on pure intuition, or on pure reason; such an idea runs counter to his scientific and empiricist predilections. Carnap appealed to the idea of analyticity as an alternative (see CARNAP). Quine, famously, also rejects Carnap's use of this idea and with it any significant idea of the a priori or of necessity. We shall discuss Quine's arguments against Carnap's notion of analyticity, or against the idea that there is a serious and significant distinction to be made between the analytic and the synthetic. Quine's rejection of the distinction, however, is only half the story. The other half is to show how he can make sense of the apparently a priori status of logic and mathematics without it – or, better, perhaps, how he can account for those facts which have led philosophers to think that logic and mathematics must be a priori.

Understanding Quine's attack on Carnap's notion of analyticity is complicated, partly because it was for a time, I think, not entirely clear even to Quine himself exactly what he is attacking and how. Quine thinks of Carnapian analyticity as truth in virtue of meaning, and so also thinks that if we had a clear understanding of the notion of meaning – more precisely, of synonymy, or sameness of meaning – then we would have gone a long way towards making clear sense of a notion of analyticity. So for a long time Quine's attack on analyticity seemed to be part and parcel of an attack on the notion of meaning, as unclear or undefined. And certainly Quine is skeptical as to how far we can make clear sense – which for him means sense in scientific, especially behavioral, terms – of the idea of synonymy. (He is not, however, skeptical of the notion of *meaningfulness*. See his essay "The Problem of Meaning in Linguistics," in Quine 1961, the burden of which is precisely that the notion of meaningfulness is not afflicted with the same sorts of problems as the notion of sameness of meaning, but lends itself to a ready, if somewhat rough, understanding in behavioral terms.) His view of synonymy is *not* that there is no sense at all to be made of it anywhere, even though he has reason to think that we may not be able to make complete sense of it everywhere. So we see Quine, as early as "Carnap and Logical Truth" (1963, written in 1954) accepting that there may be a limited notion of analyticity to be had. His willingness to accept some notion of analyticity becomes more marked as time goes by. In *Roots of Reference* (1974) he proposes a tentative definition of synonymy, and with it an understanding of analyticity; in "Two Dogmas in Retrospect" (1991) we find Quine arguing that (first-order) logic is analytic. How are we to understand this situation?

Quine continues to reject the idea of a notion of analyticity that would play anything like the central philosophical role that Carnap allotted it. In order to play that role, a notion of analyticity would have to meet two requirements. First, it would have to have the right scope: the truths of logic and mathematics, at least, must come out as analytic. Second, it must also, at least in Quine's view, mark a significant epistemo-

logical distinction: analytic truths need no justification (or else what counts as “justification” for them is wholly different in kind from the justification of synthetic truths). While Quine accepts a notion of analyticity, it is not one that satisfies either of these requirements. Let us begin with the question of scope and definition.

Quine, as we saw, takes analyticity to be truth in virtue of meaning. But how are we to understand the idea of meaning, as it occurs here? For Quine, the only thing that could be relevant to the meaning of a word or a sentence in a given language is how it is used by speakers of that language. This is an important point. Quine has been accused of being unduly behavioristic, especially about language. Certainly he has a general bias in favour of a behaviorist approach to the mind. He claims, however, that his insistence on approaching language-use behaviorally is not merely the result of prejudice. Indeed he offers an argument for some form of behaviorism in this context. The passage is worth quoting at some length:

In psychology one may or may not be a behaviorist but in linguistics one has no choice. Each of us learns his language by observing other people’s verbal behavior and having his own faltering attempts observed and reinforced or corrected by others. We depend strictly on overt behavior in observable situations. As long as our command of our language fits all external checkpoints . . . so long all is well. Our mental life between checkpoints is indifferent to our rating as a master of the language. There is nothing in linguistic meaning beyond what is to be gleaned from overt behavior in observable circumstances. (Quine 1990: 37–8)

For Quine there can be no more to meaning than is implicit in the actual use that is made of the language.

Quine’s interest is exclusively in knowledge, and the aspect of the use of language that primarily concerns him is our accepting or not accepting sentences. Thus he says, early on: “in point of *meaning* . . . a word may be said to be determined to whatever extent the truth or falsehood of its contexts is determined” (1936: 89). But then the question is: *which* of the contexts of a word must be so determined in order to determine its meaning? Without some reason to discriminate, we have no reason to treat one context as more definitive of a word’s meaning than any other. But then no true sentence in which the word appears would have any better claim to be analytic than any other such sentence; clearly no useful analytic/synthetic distinction can be erected on that basis.

What sort of thing might give us reason to discriminate among contexts? If mastery of some small subset of a word’s uses gave one mastery of its use as a whole, then there would be reason to say that those uses, those contexts, constituted its meaning. And clearly this happens in some cases. A child who otherwise has a fair degree of linguistic sophistication but does not know the word “bachelor” can be given a mastery of that word all at once, at a single stroke, by being told that bachelors are unmarried men. This fact gives us every reason to say that “bachelor” *means* “unmarried man,” and that the sentence “All bachelors are unmarried” is analytic – which Quine, at least in his later work, certainly accepts (see Quine 1991: 270). Along these lines, he proposes a definition of analyticity: “a sentence is analytic if *everybody* learns that it is true by learning its words” (1974: 79). He argues that first-order logic is analytic by this sort

of definition, but that other analytic truths will all be trivial. In particular, there is no prospect of arguing on this sort of basis that mathematics is analytic; apart from other considerations, Gödel's incompleteness theorem would be an insurmountable barrier to such an argument (see TARSKI, CHURCH, GÖDEL).

We have yet to discuss the question of the epistemological significance of the notion of analyticity. This, Quine came to see, is the crucial question; in the 1980s he wrote: "I now perceive that the philosophically important question about analyticity and the linguistic doctrine of logical truth is *not* how to explicate them; it is the question of their relevance to epistemology" (Hahn and Schilpp 1998: 207). Why should anyone think that showing a sentence to be analytic for a given language – learned in the course of learning the language – shows anything about its epistemological status? Why might one think that it shows that for that sentence no justification is required, or that the question of justification is somehow misplaced? Well, clearly it might be thought to show that *given that we are speaking that language* the question of the justification of that particular sentence does not arise. But why does the question not simply become one of the justification for speaking that language? We are presumably operating here with very tight identity-criteria for languages, so that shifting the meaning of the one word "bachelor" would mean that we were speaking a different language (if this seems excessively odd, we might speak in terms of idiolects rather than languages; but the point is the same). And given that conception of a language, it is not obviously absurd to ask for the justification for speaking a given language (it is no longer enough to say, "it is the one I was brought up with and feel most at home in," for this quality would survive minor shifts).

For Carnap, the choice of a language is in epistemologically important ways unlike the choice of a theory within a language. The former is not a matter of *correctness*, of right or wrong; it is a practical matter having to do with pragmatic factors such as the simplicity of a given language and its convenience for this or that goal. This idea issues in what he calls the "Principle of Tolerance": since choice of language, unlike the choice of a theory within a language, is not a matter of correctness or incorrectness, we should be tolerant, and allow people to work with whatever language they choose. Within a given language, justification is more or less rule-governed, governed by the rules of that particular language; justification, like other significant philosophical notions in Carnap's view, is thus language-relative. But the choice of a language itself is not something that can be justified in the same sort of way, since without a language we have no rules of justification to which to appeal.

Carnap emphasizes the distinction between the justification of choice of a theory within a language and the justification (or the lack of need for justification) of choice of language. The idea that attributing analyticity to a sentence has epistemological significance depends upon this distinction. Saying of a sentence that it is analytic would mean that it is in some sense integral to the language that we currently speak. So if we ceased to accept that sentence we would have modified the language. But that would leave open the possibility that we might have evidence which would justify that modification of the language. Thus it would seem that evidence might bear on an analytic sentence in the same sort of way in which it bears on a synthetic sentence, unless the notions of evidence and justification in play are of different kinds in the two cases. Carnap, of course, holds there there is just such a difference in kind. He claims that

there are quite different conceptions of evidence and justification at work. Within the language, justification is a rule-governed procedure, and a matter of right or wrong; when the language itself is being chosen, however, there are no rules to which to appeal, and the choice is purpose-relative and to some extent arbitrary.

Quine attacks this distinction from both sides. He denies that (internal) justification is to any significant extent a rule-governed procedure. He says, for example:

I am impressed . . . apart from prefabricated examples of black and white balls in an urn, with how baffling the problem has always been of arriving at any explicit theory of the empirical confirmation of a synthetic statement. (1961: 41–2)

He also insists that all our cognitive choices, including the choice of a language for knowledge, are directed towards the same end: achieving the most successful theory, where a crucial test of success is the generation of true predictions. Vaguer virtues, such as simplicity and fruitfulness are also relevant. These are the sorts of things that Carnap counted as “pragmatic factors,” applicable to questions of language-choice. Quine claims that they are applicable also to what Carnap would count as empirical beliefs. They may not in any very obvious way be applicable to the question whether there is now a desk in front of me, but certainly they are to more or less abstract claims of theoretical physics. For Quine there is a continuum here, with no sharp breaks to be had.

Quine thus holds that even where we have a significant truth which is analytic, this status simply does not matter epistemologically:

“Momentum is proportional to velocity” counts as analytic. But do we care? Einstein’s relativity theory denies the proportionality law, complicating it with a formula involving the speed of light. But instead of accusing Einstein of a contradiction in terms, we simply stand corrected. (Quine 2000)

Now Carnap might agree that we “stand corrected” because we accept that Einstein has shown us that a non-Newtonian language works better for making some predictions, but he would insist that this is a different sense of correction from that in which we are corrected when we change our mind about a belief that does not involve a change of language. But this is precisely what Quine denies, as we have seen.

One issue which arose above was Quine’s view of the nature of justification, the relation between the evidence we have and the beliefs that we hold on the basis of it. The point there was that there is not, in general, a simple relationship between a sentence, on the one hand, and an observation or group of observations that justify it, on the other hand. Justification is not, in general, a simple and rule-governed matter. Of course there are sentences, such as “there is a desk in front of me now,” which do seem to have a very straightforward relation to observations. What makes that case straightforward is that it hardly matters what else a person believes: given the right observations, almost anyone will accept that there is currently a desk in front of them. The justification relation here holds between observations and the individual sentence believed, whatever one’s other beliefs may be. In Quine’s view, however, this is a poor paradigm to use for knowledge as a whole. In general, justification is *holistic*, meaning that it does not apply

to sentences taken individually, in isolation from others, but rather to larger or smaller chunks of theory, made up, in some cases, of a large number of sentences. Many of the sentences we accept – most obviously the more abstract and theoretical ones – have relations to observations only if we tacitly assume many other sentences. These other sentences, background assumptions, are required if the sentence in which we are interested is to have any observational consequences at all. From the point of view of the working scientist, the background assumptions may be confidently accepted, and only the individual sentence up for testing. From a more abstract point of view, such as Quine's, however, what is tested by observation is not the individual sentence alone, but rather the whole set of sentences that implies the observational consequences. From a sufficiently abstract point of view, indeed, it is always the whole of our knowledge that is tested. Any test of a sentence presupposes truths of logic among its background assumptions. Logic, however, is used everywhere in our system of beliefs, so in a rather Pickwickian sense it is that system as a whole that is at stake. (Quine calls this extreme holism "legalistic" (1991: 268). I take this to mean that it holds from a very abstract point of view, but not that it is unimportant.)

Holism is not new with Quine. When "Two Dogmas" was reprinted, Quine added a footnote to Duhem, and there is a clear statement of the view in Carnap's *Logical Syntax of Language*, with references to Poincaré as well as to Duhem (Carnap 1937: 318). Quine's uses of the doctrine, however, are novel. One use we have seen: it is the basis of the claim that justification (within a language) is not the sort of rule-governed procedure that Carnap sometimes suggests, and so is not different in kind from the sort of justification that applies to the choice of one language rather than another. This claim, in turn, is crucial for Quine's view of analyticity and the a priori, discussed in the previous few pages.

A second use that Quine makes of holism also relates to the question of the a priori, but in a quite different way. He attempts to undercut the idea that there must be a priori knowledge by invoking holism to explain the phenomena which led some philosophers to invoke the idea of the a priori, but to do so without invoking that idea. (If those phenomena can indeed be explained without the a priori, then their existence no longer constitutes a reason to accept the a priori.) There is no doubt that the theorems of mathematics and of logic are not discovered by experiment or, at least in any ordinary sense, justified by observation. And their falsehood seems completely inconceivable. How can these facts be explained in accordance with Quine's views? Holism provides the answer. Logic figures everywhere in our system of beliefs; mathematics is used in many branches of knowledge. No one observation or experiment bears on them, but the success of our system of beliefs as a whole in predicting experience provides justification; this justification is exceedingly indirect. For Quine that only puts it at one end of a continuum which we already have reason to accept, for " $e = mc^2$ " is justified very much less directly than is "there is a desk in front of me now." Equally the unimaginability of the falsity of logic becomes intelligible. Given the ubiquity of logic, changing it means making changes everywhere in our system of knowledge. It is not to be wondered at if this is hard to conceive.

Our discussion so far has been focused on Quine's rejection of the a priori. He is thus left with no kind of knowledge other than the ordinary knowledge of common sense

and (better) science. Philosophy too, since it claims to yield knowledge, must be of this same general sort. This is the doctrine that Quine calls naturalism; it is absolutely fundamental to his thought. The rest of this essay will take naturalism as its starting point, and investigate Quine's philosophy as an unfolding of that doctrine. I shall begin with topics having to do with epistemology, and then move to topics whose focus is ontology and metaphysics (or its Quinean analogue). There is, however, one further issue which we should briefly mention here.

Quine has argued for the possibility that two translators might come up with different translations of some sentences – not merely stylistically different, but “not equivalent in any plausible sense of equivalence, however loose” (Quine 1960: 27). This is the controversial doctrine known as the indeterminacy of translation. Where sentences are not asserted or denied on any very direct observational basis, the evidence (behavioral evidence, of course) for one translation over another is mediated by other sentences; alternative translations of all of them might cancel out, leaving each of two overall schemes of translation as equally justified. In Quine's view there would, in such a case, be no right and wrong, no fact of the matter: it would not be our knowledge which was lacking, but rather that there was no fact to be known. This idea provoked an enormous amount of discussion in the 1960s and early 1970s, and some commentators have even thought that it is what really underlies Quine's objections to analyticity. In my view, however, it is not of great importance to his thought, except as dramatizing the idea that meaning must ultimately be answerable to behavior (see Hylton 1990). Note in this connection that Quine now speaks of indeterminacy as “a conjecture” (see Hahn and Schilpp 1986: 728).

Knowledge and the realm of the cognitive

How should we conceive of knowledge, if we are to take a scientific approach to it? Fundamental to Quine's thought is the idea that knowledge is to be understood as a biological phenomenon. Human knowledge is thought of as a condition of the human animal. It originates in the struggle of one species of primate to survive. The opening sentences of *From Stimulus to Science* read like this: “We and other animals notice what goes on around us. This helps us by suggesting what we might expect and even prevent, and thus fosters survival” (Quine 1995: 1). This is what one might call a Darwinian conception of knowledge: knowledge as an adaptive mechanism, fostering the survival of the species.

In this view we see a decisive rejection of the long-standing philosophical tradition that sharply distinguishes between real knowledge and mere belief or opinion, between *scientia* and *doxa*. That tradition tends to assimilate real knowledge – *scientia* – to knowledge of a proposition of mathematics, when it is known on the basis of a thoroughly understood proof. Real knowledge is accordingly thought of as infallible and known with certainty. None of these ideas fits with the idea of knowledge as a biological phenomenon, as what helps the human animal to get by in its dealings with the world. Some philosophers have held that those ideas are implicit, more or less, in the word “knowledge.” Quine's response is simply to abandon that word for “scientific and philosophical purposes” (1987: 109). He continues to use it informally (as shall I in expounding his views), but without the weight that it may be thought to carry in

general use. He makes no sharp distinction between our knowledge, and that which we, or experts among us, accept upon reflection. Quine's conception of cognition, then, is fallibilist through and through: no part of our system of beliefs can be counted as wholly immune from revision, though some parts are no doubt far more secure than others.

A second very general point about Quine's conception of knowledge is that he takes it to be linguistic or verbal; at least for "scientific and philosophical purposes," he thinks of our system of beliefs as being embodied in sentences:

What sort of thing is a scientific theory? It is an idea, one might naturally say, or a complex of ideas. But the most practical way of coming to grips with ideas, and usually the only way, is by way of the words that express them. What to look for in the way of theories, then, are the sentences that express them. (Quine 1981: 24)

Much of Quine's interest in language and in its analysis arises from the fact that our knowledge is embodied in language.

In one way, as we saw, the idea of knowledge, as Quine employs it, marks no very sharp distinction. Unlike some philosophers, he does not use the term as an honorific, connoting some particular high degree of justification or of certainty. He does, however, make a sharp distinction between the realm of the cognitive and the rest of human activity. To call a human activity cognitive is, roughly, to say that it is answerable to, if not exclusively aiming at, predictions of sensory experience. It is perhaps a presupposition here that cognitive activity and cognitive language can be peeled off from the chaotic mass of human activity and language generally – or at least that we can abstract without distortion, and talk of the cognitive while ignoring the rest.

The prediction of experience is a practical matter; at the limit, as we indicated, survival is at stake. This is the sense in which Quine's conception of knowledge is Darwinian. Thus far Quine is with the pragmatists. On the other hand, tying the concept of knowledge to the prediction of sense-experience enables Quine to make clear distinctions, and to erect barriers, in places where the pragmatists would not. Activities which predict experience are cognitive; others, though they may contribute to human flourishing in other ways, are not. The justification for this distinction is presumably (for Quine is not explicit here) that sense-experience is our only way of finding out about the world (we shall discuss this idea shortly). In spite of Quine's practical, Darwinian view of knowledge, there is thus a sense in which his view makes a clear distinction between the theoretical and the practical. *Theoretical* success is success in prediction of sensory experience.

Evidence

As traditionally conceived by philosophers, this notion includes two strands. On the one hand, evidence is thought of as consisting of the epistemologically most fundamental items of our knowledge: evidence is that which is, so to speak, first in the order of knowledge (that which is evident), and so also that from which other items of knowledge must be inferred. On the other hand, evidence is also to consist of immediately given data, devoid of any conceptual impositions of our own; since no interpretation is

involved, there is no room for doubt. These two strands are in tension. What is literally first in the order of knowledge seem to be facts about other people and ordinary physical objects. Yet sentences recording such facts do not seem simply to record raw data. They involve conceptualization, and are (notoriously) open to doubt.

Quine's response to these difficulties is to abandon the traditional conception of evidence completely, in favour of a physicalistic alternative. He speaks not of "the given" but of the stimulation of our sensory receptors, and of observation sentences: roughly, sentences that any speaker of the language is disposed to accept or reject simply on the basis of current stimulation. Thus he says:

our immediate input from the external world [is] the triggering of our sensory receptors. I have cut through all this [i.e. the difficulties of analyzing the notions of observation and experience] by settling for the triggering or stimulation itself and hence speaking, oddly perhaps, of the prediction of stimulation . . .

Observation drops out as a technical notion. So does evidence, if that was observation. We can deal with the question of evidence for science without the help of "evidence" as a technical term. We can make do instead with the notion of observation sentence. (1990: 2)

What is at stake in philosophical talk of evidence is, from Quine's point of view, the issue of how we find out about the world. Here we have a crucial example of his method, of the idea that the study of knowledge is to be naturalized, and as far as possible put on a physicalistic basis. For him, the issue is to be taken as a scientific question: how do we come by information about the world? And the answer is that we do so by the impact of various forms of energy on our sensory surfaces. Physics will tell us what forms of energy there are; physiology and psychology will say which forms of energy human beings can detect, i.e. to which forms human beings respond. The central fact here is that it is only through stimulation of our nerve-endings by energy impinging on our sensory surfaces that we human beings know anything at all about the world. This is fundamental to Quine's epistemology and it is, he emphasizes, "a finding of natural science itself" (1990: 19).

Quine's use of the notion of stimulations is thus symptomatic of his general shift in perspective. The question is not: what is given to me, at the outset of my cognitive endeavors? But rather: how do we humans gain knowledge of the world? This is on his view a straightforward scientific and causal question. The answer refers us to stimulations of our sensory surfaces. Quine thus abandons, or greatly modifies, the traditional concept of sensory evidence. In so doing he shifts the question to which the original conception was an answer. The traditional philosopher's demands are not met by Quine's view. In the most straightforward sense, the occurrence of stimulations is independent of and prior to theory. Our *knowledge* of such matters, however, is clearly not independent of theory. So there is here no prospect of what the traditional philosopher sought: support of the theory which is wholly independent of the theory. To the extent that this strikes us a problem or a paradox, to that extent we have not accepted Quinean naturalism.

The point we have just noted links Quine's rejection of the sensory given with his rejection of the traditional conception of the a priori, or of any analogue. The given

and the a priori were each conceived, by some philosophers, as a kind of extra-theoretical knowledge, knowledge somehow free of the vicissitudes affecting the ordinary knowledge of common sense and science. Quine denies that there is any such kind of knowledge. Within what we take ourselves to know there is, no doubt, better and worse; there is, however, no knowledge of a wholly different and superior kind. Quine is consistent here. He does not take even the most fundamental points of his own philosophy to be a priori – and this includes the doctrine that we know about the external world through impacts on our sensory surfaces. While it is extremely unlikely, there are imaginable circumstances under which we might drop our present idea of sensory evidence entirely. If sufficient confusion resulted from our following the evidence of our senses, as we now understand them, and sufficient success on the part of those who claim to hear voices in their heads, say, our estimate of the role of stimulation of our sensory surfaces might change. (This point is explicit; see Quine 1990: 20f.)

For Quine, as we have said, the central fact about knowledge is that it is only through stimulation of our nerve-endings that we know anything at all about the world. Such stimulations provide the only empirical constraint on our system of beliefs, the only external criterion of success. (By speaking of an *external* criterion I mean to leave room for what one might think of as internal factors: such as the overall simplicity of the system.) This fact suggests that there may be empirical slack between evidence, even the totality of all possible evidence, and theory. We can focus this idea by asking whether there might be two systems of belief, different from one another but each fully successful at “predicting stimulations.” Quine’s answer, though somewhat qualified and complicated, is that nothing rules this out. This is the doctrine known as *the underdetermination of theory by evidence*. Even if, *per impossibile*, all the evidence were in, still any given theory based on that evidence might not be uniquely justified. If we had a theory that explained and predicted the evidence satisfactorily, that would of course justify it; since it would at least in principle be possible for another theory to do as well, however, the justification would not be unique.

The relation of evidence to knowledge: observation sentences

So far we have said something about evidence and something about knowledge, as Quine conceives them, but nothing explicit about the relation between the two. Clearly there is a gulf between the sensory stimulations that are our only source of information about the world, and the mass of sentences in which our beliefs about the world are embodied. How is this gulf bridged? For the sentences most directly tied to sensory evidence, *observation sentences*, there is a fairly clear account. The other sentences of our theory of the world get their relation to evidence *via* their relation to observation sentences, and here the account is much sketchier. (A full account is, Quine thinks, not yet available, and may never be.) Discussing observation sentences will at least give us an idea of how the gap between evidence and theory is bridged.

The first point to make about what Quine calls observation sentences is that they are – by his lights – *sentences*. This does not mean that they are all sentences in the grammatical sense: “Rabbit,” taken as a complete utterance, is Quine’s own example. The idea of a sentence here is that of a piece of language which may be used to

say something, and thus may be true or false. For Quine, the fundamental evidential relation holds between sensory stimulations and units of language of that sort.

We can bring out the significance of this idea by contrasting sentences (as we are using that word) with referring expressions, or “terms,” as Quine says. In most uses, the word “rabbit” is a term, not a sentence: if I say “I see a rabbit” then in that sentence it is a term which functions as part, but only part, of a sentence. Mastery of the word in that sort of use, Quine claims, requires more than a mere ability to respond to the presence of rabbits. It requires also that one can distinguish the circumstances that license the claim “There’s one rabbit” from those licensing the claim “There are many rabbits”; that one can similarly distinguish “There’s the same rabbit again” from “There’s another rabbit”; and so on. The ability to make these and related distinctions requires some knowledge of the ways of physical objects in general, and of rabbits in particular. Sentences containing terms are thus not epistemologically basic. Mastery of the use of a term already requires that one possess some knowledge, and then the question of the evidence for that knowledge must arise. Mastery of the word “rabbit” as an observation sentence, by contrast, requires only the capacity to respond to environments containing rabbits in ways in which one does not respond to environments which do not contain rabbits; this capacity requires no auxiliary knowledge. This kind of ability is all that is presupposed by the mastery of an observation sentence. It is primitive and fundamental. It is what underlies all cognitive language-use, including the most sophisticated, but other forms of language go beyond it in principle.

To this point we have said little more about observation sentences than that they are sentences, though that idea has proved to be far from trivial. Beyond that, an observation sentence is what Quine calls an “occasion sentence,” i.e. it is one that is true when uttered under some circumstances and false under others (so “It’s raining in London” is an occasion sentence; “Gold is a metal” is not). The rough idea that Quine intends to capture is that an observation sentence is an occasion sentence about which there is community-wide agreement under any given circumstances. A little more precisely, we can distinguish two conditions which an occasion sentence must satisfy to count as an observation sentence. First, whether a given speaker of the language is disposed to assent to, or dissent from, an observation sentence at a given time is simply a matter of the stimulations that that individual is undergoing at that time. For each individual, the same stimulation-pattern will typically lead to the same verdict each time. Second, any fully competent speakers of the same language, in the same circumstances, will agree on an observation sentence; Quine speaks of “unhesitating concurrence by all qualified witnesses” (1995: 44). (There is some vagueness here, arising from the vagueness of “in the same circumstances,” and of “qualified witness.”)

Observation sentences thus assert the presence of something readily detectable by the senses. Whether such a sentence is correctly assertable in a given situation does not depend upon ancillary information, unless it is shared by all speakers of the language. Hence “It’s cold here!” might qualify, but “That’s Quine!” would not, since not all English-speakers would recognize that philosopher on sight. The range of observation sentences will vary with our decision as to exactly who should be included among the fully-functioning speakers of the language. “That’s red” will presumably count if we exclude the blind and the color-blind, but otherwise not. Observation sentences are the epistemologically most basic parts of our theory of the world. They can be known before

anything else. (They thus play one of the roles traditionally accorded to the notion of evidence: they are first in the order of knowledge.) The knowledge that they embody is so rudimentary that in almost all cases it goes unspoken, but can be elicited by raising the sentence as a question and noting the subject's reaction.

Considered holophrastically, i.e. as unanalyzed wholes, observation sentences are simply responses to stimulation, and are in only the most minimal sense conceptual or theoretical. Hence such sentences can be mastered by a child otherwise quite innocent of language. As Quine says, "Their direct association with current stimulation is essential if the child is to acquire them without prior language" (1990: 5). Considered as made up of parts, however, they connect with sophisticated theory, for the words which make them up recur in more theoretical contexts. This dual aspect is essential to the function of observation sentences as the starting point of language and conceptualization. Because they presuppose so little, observation sentences will be the first sentences learned by a child (or, indeed, by an adult trying to find his or her unaided way in a wholly strange linguistic community); their learning presupposes no prior conceptual or theoretical resources. Because their terms recur in higher theory, learning such a sentence is a start on learning the language as a whole; it is only this sharing of vocabulary which unites the observation sentences with the rest of the language (1990: 8).

Observation sentences, we saw, are occasion sentences: the truth-value of such a sentence will vary from one occasion of utterance to another. Our scientific theories, however, and most other serious knowledge, consists of standing sentences: sentences true or false once-for-all. How is this gap in turn to be bridged? Quine's answer appeals to the notion of an *observation categorical*. This is a sentence compounded of two observation sentences, saying that whenever one of them holds the other will also hold: "Whenever there is smoke there will be fire," for example. This is a standing sentence, and so might be implied by a serious branch of organized knowledge. On the other hand, both of its component parts are observational (or so we are supposing). So we can tell right off whether a given situation is one in which there is smoke, and whether it is one in which there is fire. Hence we can tell right off whether a given situation is one in which the observation categorical is falsified. Because the sentence is in effect a generalization over all situations, it cannot, of course, be *verified* by a single situation, but it – and hence the theory which implies it – can be falsified. Quine readily accepts this asymmetry between verification and falsification, which fits with his general fallibilism. Since our theories have infinitely many observational consequences, they cannot be conclusively verified; in principle, however, a single observation may falsify a theory.

Naturalized epistemology and normativity

We have been articulating Quine's general conception of knowledge, evidence, and the relation between them. This conception is thoroughly naturalistic: Quine treats knowledge as a natural phenomenon, to be studied by the procedures of science. Most of Quine's own work in epistemology is an articulation and defense of this very general conception. He suggests, however, that there is also room in epistemology for detailed piecemeal work of a more recognizably scientific kind. This would consist in tracing the

connections between theory and evidence in a psychologically realistic fashion, to see how our knowledge is in fact related to the evidence that we have. Epistemology of this sort is thus a branch of psychology. (Of course psychology is itself among the items of knowledge whose relation to evidence is to be investigated in this way; Quine speaks here of “reciprocal containment” (1969: 83).)

Epistemology, as traditionally thought of, is a normative subject: it aims to tell us not merely about what is but also about what ought to be; it aims to tell us not only what we do in fact believe, and on what evidence, but also which beliefs are justified on the basis of the evidence that we have. Quinean epistemology, at least according to his account of the matter, is descriptive. To what extent, if any, can this descriptive subject take on the burden of traditional normative epistemology? This is a large and complex question (see Gregory 1999, to which I am indebted here). Roughly we may say that Quine has no room for the very large-scale questions and doubts which are one kind of starting point for traditional epistemology. He has no sympathy at all, for example, with global skepticism. The aim of our knowledge is to predict sense-experience. A theory that does that satisfactorily does all that we can ask. There is no further question as to whether it tells us about the nature of reality:

what if . . . we have achieved a theory that is conformable to every possible observation, past and future? In what sense could the world then be said to deviate from what the theory claims? Clearly in none. (Quine 1981: 22)

We cannot divorce the idea of reality from that sense-experience. (This point will emerge further in the next section, below.)

There is, nevertheless, a sense in which Quinean epistemology is normative. The criterion of success for all putative knowledge, for science in Quine’s broad sense, is the prediction of sense-experience. Quine sees this as defining the notion of science (1990: 20). This definition is not arbitrary: our primary aim, in science, is to find out about the world, and one thing we know – a well-established piece of scientific knowledge – is that it is only through sense-experience that we come to know about the world. Given that our goal is fixed, there are questions of a normative sort about the best ways in which to achieve that aim, and in this instrumental sense epistemology is normative:

Naturalization of epistemology does not jettison the normative and settle for the indiscriminate description of ongoing procedures. For me normative epistemology is a branch of engineering. . . . There is no question here of ultimate value, as in morals; it is a question of efficacy for an ulterior end, truth or prediction. The normative here, as elsewhere in engineering, becomes descriptive when the terminal parameter is expressed. (Hahn and Schilpp 1998: 664–5)

Realism

As we have seen, Quine’s naturalism can be identified with the view that there is essentially only one kind of knowledge. In particular, there is no special philosophical perspective from which we can attain knowledge that is independent of our

ordinary scientific or commonsensical theory of the world. Thus on his account we are always inside that theory, modifying it, perhaps, but not wholly transcending it. There is no transcendental standpoint that is independent of our ordinary knowledge, and from which we can evaluate that knowledge without presupposing it. Quine's realism is an important application, and illustration, of this view.

Quine's work is full of remarks which might suggest that he does not take our theories – including the “theory” that is common sense knowledge – to be (really) true, and does not take the objects that those theories presuppose to be (really) real. He says that our theories far outrun the evidence that we have for them, and that more than one theory is compatible with that evidence. The disparity between our evidence and our knowledge is a recurrent theme in his work. His insistence on this disparity, together with his view that our knowledge is justified by its efficacy in predicting and understanding the course of experience, might lead one to suppose that his is an instrumentalist or pragmatist view: that our theories, even at their best, are not really *true*, that they do not aim to correspond to an extra-theoretical world, but are simply useful instruments for predicting, understanding, and controlling future experience. Some critics have taken this to be Quine's view, and have seen his insistence on realism as a contradiction, or as a mis-statement of his actual position (see, for example, Lee 1986, and Smart 1969). Quine, however, insists that his view is “robust realism” (1990: 21), and that appearances to the contrary can be dispelled by taking naturalism seriously enough.

On Quine's view, the objects that our theories presuppose *do* exist in extra-theoretic reality. It is part of our theory – that is, part of the best understanding of the world we have – that those objects (with a few exceptions, most of them straightforward) are not dependent on us or our theorizing. Now the critic may protest: it may be part of our theory that our theoretical objects really exist, independent of our theory – but that's just part of the theory. Do the objects *really* exist? But this is an attempt to ask a question from a stance independent of our theorizing about the world: the point of Quine's naturalism is that there is no such stance. The objects that we believe in exist, and are real, in the only sense of those ideas that we actually have.

More generally, the apparently skeptical remarks that Quine makes when discussing our acquisition of knowledge do not affect his belief in the truth of the knowledge that is thus acquired: his ontology is in this respect insulated from his epistemology. This is a consequence of Quine's version of naturalism, and in particular of the reciprocal containment of science, with its ontological claims, within epistemology, and vice versa. Let us see how this goes. We accept, let us suppose, the best overall theory of the world that is available. Freely drawing on this theory we do epistemology, i.e. we investigate the way in which human beings – including ourselves – come to formulate theories and to posit the existence of objects; among the theories thus investigated are those that we are drawing on in the course of our investigation. Now the crucial point is that to call a body of knowledge a “theory,” or to call an object a “posit,” does not in the least impugn its truth or its reality. “Theory” here is not mere theory, contrasted with real knowledge, for *any* body of knowledge will count as a theory from the point of view of epistemology; nor is “posited object” contrasted with real object. Thus, as Quine famously says:

To call a posit a posit is not to patronize it . . . Everything to which we concede existence is a posit from the standpoint of a description of the theory-building process, and simultaneously real from the standpoint of the theory that is being built. *Nor let us look down on the standpoint of the theory as make-believe; for we can never do better than occupy the standpoint of some theory or other, the best we can muster at the time.* (1960: 22; my emphasis).

The crucial point is the one emphasized: there is no alternative to occupying some substantive theory of the world, and to do this means accepting that theory, at least for the moment, as true, and accepting its objects as real. Of course we may develop our theory into a different one, but we cannot occupy some neutral philosophical vantage point; nor can we accept a theory while still pretending that we are not accepting it as true.

In Quine's use, then, theoretical knowledge is not contrasted with ordinary knowledge. Similarly, theory for Quine is not contrasted with fact. All knowledge is theoretical, in Quine's sense. Just as there is here a stretching (or a distortion) of the word "theory," so also there is a stretching of the word "posit." When we say of, neutrinos, for example, that they are posits, we would generally be taken to mean that some person or group of people consciously posited them. Quine, however, speaks of physical objects in general as posits, and here there is no such implication. No conscious decision was ever taken to posit such things. As in the case of "theory," Quine's use thus assimilates ideas which one might suppose to be importantly different.

Metaphysics and regimentation: logic and extensionality

To this point we have seen Quine reflecting on the nature of our knowledge, its sources and bases, and on its status. The philosopher's task, as he conceives it, also includes clarifying our knowledge and helping us to attain a clear view of just what it comes to and what it really commits us to. This latter kind of task may be thought of as the Quinean version or analogue of metaphysics. Like the others, it aims, in Quine's view, to contribute to the overall scientific enterprise. It should, for example, enable us to avoid useless or misleading questions, help to suggest fruitful lines of further inquiry, and expose potential problems that may lurk in scientific theories.

How is this task to be approached? The method here is to show how various parts of our knowledge could be reformulated in the clearest possible terms. We have reason to take the objects and categories revealed by this reformulation as real and fundamental. An extended passage from *Word and Object* is worth quoting at length on this topic:

The same motives that impel scientists to seek ever simpler and clearer theories adequate to the subject matter of their special sciences are motives for simplification and clarification of the broader framework shared by all the sciences. Here the objective is called philosophical, because of the breadth of the framework concerned; but the motivation is the same. The quest of a simplest, clearest overall pattern of canonical notation is not to be distinguished from a quest of ultimate categories, a limning of the most general traits of reality. Nor let it be retorted that such constructions are conventional affairs not dictated by reality; for may not the same be said of a physical theory. True, such is the nature of reality that one physical theory will get us around better than another; but similarly for canonical notation. (1960: 161)

Themes that we examined in the first part of this essay, having to do especially with the rejection of the Principle of Tolerance and with the remoteness of some parts of our theory from experience, re-emerge here. Simplicity, clarity, and convenience are not merely “pragmatic” virtues which are to be distinguished from the cognitive or theoretical virtue of truth. If we have a language which displays those virtues to the highest extent, then we have reason to take the structure of that language as telling us something about the real world, just as we take theories which display those virtues to tell us something about the world.

One aspect of the philosopher’s task, then, is to find the clearest and simplest framework in which to formulate our knowledge. In Quine’s view, this clarity and simplicity is to be found in first-order logic. (By “first-order” logic is meant logic of truth-functions and quantifiers which bind variables in positions occupied by singular terms, but not positions occupied by predicates.) He takes it that our knowledge is at its clearest when it is formulated in the syntax of first-order logic: a syntax which uses only truth-functions, predicates, variables, and quantifiers. This is not to say that he advocates language reform, or that he thinks that we should in fact reformulate all of our knowledge in those terms. But where our concern is with getting clear about what some part of our knowledge really commits us to, we would do well to consider how it might be phrased in logical syntax. This syntax is extraordinarily transparent and economical, which makes theorizing about it simple, and yet has surprising expressive power. It contains neither proper names nor function-symbols among its primitive expressions, for example, yet the effect of each can be easily achieved. (The basic techniques of doing so derive from Russell’s Theory of Descriptions (see RUSSELL).) One moral to be drawn here is that we may be able to achieve the effect of a particular construction without in fact having to augment our stock of primitives. Quine is much concerned to take advantage of this kind of economy wherever it is possible, for the sake of the clarity and simplicity of the overall theory. (In ontology, as we shall see, Quine has a similar concern with economy: to show how particular kinds of object which appear to be assumed in what we take ourselves to know need not in fact be taken for granted as primitive, because their effects can be duplicated by other means.)

One feature of this framework which has proved extremely controversial is its *extensionality*. A language is extensional when any sentence of it retains its truth-value under any one (or more) of three kinds of changes: (1) any name in it may be replaced by any other name of the same object; (2) any predicate in it may be replaced by a co-extensive predicate (i.e. one true and false of exactly the same objects); (3) any sentence embedded in another sentence may be replaced by a sentence of the same truth-value. It is important to note that these three requirements which a language must satisfy to count as extensional interlock in ways that make it very hard for a reasonably comprehensive language to satisfy any one of them without satisfying all three. This point is important because the first requirement is, on the face of it, far more plausible than the third. It is extremely plausible that if we really understand what a given sentence is about, then we could replace the name of that object in the sentence with another name of the same object without altering the truth-value of the original. If a sentence is genuinely *about* an object, how we name that object should be a matter of indifference to the truth or falsehood of the sentence, though it may affect, for example, its poetic quality. (One might, indeed, take this idea as a partial definition of the somewhat

vague idea of *aboutness*.) No such superficial plausibility attaches to the third requirement, yet it turns out to be very hard to see how one can accept the first without also accepting the third. (See, for example, Quine 1995: 91–2.)

Any language which uses the syntax of first-order logic (along with any standard semantics) is, in virtue of that fact, an extensional language. Yet it would be a mistake to think Quine accepts extensionality simply in order to be able to use the syntax of logic. On the contrary: he takes one of the advantages of that syntax to be that it enforces extensionality, which he holds to be desirable for its own sake:

I find extensionality necessary . . . though not sufficient, for my full understanding of a theory. In particular, it is an affront to common sense to see a true sentence go false when a singular term in it is supplanted by another that names the same thing. What is true of a thing is true of it, surely, under any name. (1995: 90–1)

Quine's insistence on extensionality is a very long-running theme of his thought, going back as far as the clarification of *Principia Mathematica* (Whitehead and Russell 1910–13) in his doctoral dissertation. (See the first paragraph of this essay; also Quine 1991: 265–6.)

The requirement of extensionality has been controversial because there are large areas of discourse which, at least if taken at face value, are not extensional. This for Quine is reason enough not to take such discourse at face value. Instead, he thinks, we should either try to reformulate it so that it becomes extensional or else exclude it from the more scientific and respectable parts of our knowledge, those parts which we take as really telling us about the objective world. We shall discuss three examples; our discussion of each will be very brief, although the second and third of them are issues which have generated much controversy, and on which Quine has, largely for that reason, written extensively.

The first sort of prima-facie violation of extensionality is almost trivial; indeed it might be said that the idea that we really have such a violation here is simply a mistake. This is the case of quotation, which is worth examining because it functions as something of a paradigm for Quine. It is true to say: “‘Quine’ has one syllable”; it is false to say “‘The author of *Word and Object*’ has one syllable”; yet the one sentence might seem to be obtained from the other by replacing a singular term with another singular term designating the same object, since Quine, of course, *is* the author of *Word and Object*. Here the solution to the apparent puzzle is easy. We should not construe the subjects of the sentences as referring to a person, but rather to the words, the expressions themselves. It is of the *word* “Quine,” not of the philosopher Quine, that we say it has one syllable. And then of course the substitution no longer replaces a name by another name for the same object, since the two *expressions* are not the same.

The second sort of example is far less easy to dismiss. The case of indirect discourse, where one reports the speech or the thoughts of another, also gives rise to prima-facie cases of non-extensionality. Othello (supposing him for the moment to be a real person) has the false belief that Desdemona loves Cassio, but does not (presumably) have the equally false belief that the moon is made of green cheese. Clearly, in statements of the form “A believes that *p*” it is not only the truth-value of *p* that is relevant to the truth-value of the whole. Similarly in the case of singular terms. To use an example of

Quine's, Tom may believe that Cicero denounced Catiline without believing that Tully denounced Catiline, for he may not know that Cicero is Tully.

One response to such cases is to say that what is believed or disbelieved is not a fact about a person (or other object) but rather a *proposition*. (For this reason, philosophers often speak of such cases as statements of "propositional attitudes.") Then it is claimed that we have two distinct propositions: first, that Cicero denounced Catiline, and, second, that Tully denounced Catiline. So construed, belief-contexts become extensional, for we can no longer obtain a falsehood from a truth by substituting a co-designative expression (since this now means an expression referring to the same proposition). Quine rejects this idea, chiefly on the grounds that no clear and precise identity-conditions have been given for propositions, which should therefore not be accepted for scientific and philosophical purposes. This is of a piece with his having some degree of skepticism about meaning, for one can think of a proposition as being the meaning of a declarative sentence. (In that case the question of the identity-conditions of propositions is the same as the issue of synonymy for declarative sentences.)

Quine's treatment of indirect discourse relies not on the meanings of sentences but simply on the sentences themselves. Formally, he assimilates them to cases of quotation, by construing belief (and doubt, and hope, and so on) as attitudes towards *sentences*. Perhaps there is something odd or counterintuitive about speaking of believing or not believing a sentence, but this kind of oddity is something that Quine is, as we have seen, fully prepared to accept in the interests of clarity (as he conceives it). And if we do talk this way, then there is nothing puzzling in thinking that Tom, to revert to our example, may believe the one sentence without believing the other, for they are distinct objects.

Two points call for comment. First, Quine was for a period convinced that there are two kinds of belief. There is the ordinary kind, variously construed as an attitude towards a sentence or a proposition. But then there is also, he thought, another kind, *de re* belief, which is belief genuinely about an object, so that the way in which the object is described is irrelevant to the truth of the belief-attribution. He spent much effort trying to make sense of cases which appear to be of this sort, but subsequently came to think (quite correctly, in my view) that the appearance of two distinct kinds of belief is mistaken, and that the phenomena that he was attempting to understand are in fact not really cases of a different kind of belief. Second, it would be a mistake to think of the difficulties of making clear sense of belief contexts as merely formal. It is Quine's view, after all, that we want an extensional language not for its own sake alone, but because lack of extensionality is a sign of lack of clarity. His re-construal of belief as an attitude towards sentences puts the emphasis in what he thinks, for independent reasons, in the right place. If we are construing evidence as austere and strictly as possible, the evidence for a belief-ascription is simply a report of what the person concerned said – the very words uttered – and did. In reports of these sorts there is no violation of extensionality. Beyond that, belief-ascription relies on empathetic projection: the ascriber imagines what it would be like to be in the believer's situation. Here, it may be thought, we are out beyond the realm of hard fact.

A third area of discourse in which extensionality fails, at least on Quine's view, is that concerning modality. Suppose it said that nine is *necessarily* greater than five (the example, again, is Quine's). Replacing "nine" by a co-designative expression we can

obtain: the number of the planets is necessarily greater than five. Quine, however, sees no hope of making sense of any notion of necessity according to which this is true (cf. MARCUS). He distinguishes a notion of necessity of this sort from one which attaches to sentences (and hence can be assimilated to the case of quotation) or treats it as a statement operator. In these cases the necessity is supposed to hold of the sentence, not of the object, so there is no risk of the sort of violation of extensionality indicated above. The general issue of modality, however, is not a crucial one from a Quinean point of view, for he sees no need to accept any notions of necessity or possibility. He holds that a reconstruction of our knowledge, or of those aspects of it which we want to take with full seriousness, as telling us about the objective world, require no such notions.

We have briefly discussed three cases of prima-facie non-extensionality. It is worth emphasizing, however, that these three do not exhaust the matter. Counterfactual conditionals, statements of causality, and, if taken at face value, statements of time and tense, are among other such cases.

Ontology and its relativity

A crucial part of the task of getting clear about what our theories commit us to, on Quine's account, is gauging their *ontological commitments*. By the ontological commitments of a theory Quine means what entities that theory says there are in the world. In Quine's view, the way to settle this is by seeing what objects must be in the range of the theory's variables if it is to be true. Quine's emphasis on ontology, and on the range of variables as the measure of ontological commitment, is perhaps in part to be explained by his early work in set-theory, and his abiding interest in that subject. In that context, the range of the variables of the theory is a natural measure of its strength, and the threat of paradox makes this matter of vital concern. (The importance of this point was emphasized to me by Stephen Menn, to whom I am grateful.)

Quine takes ontology to be a product of self-conscious scientific and philosophical reflection. He does not see the philosopher's ontological task as that of capturing in perspicuous form the ontology implicit in ordinary thought and discourse, for he insists that there is no such "ordinary ontology":

a fenced ontology is just not implicit in ordinary language. The idea of a boundary between being and non-being is a philosophical idea, an idea of technical science in a broad sense. Scientists and philosophers seek a comprehensive system of the world, and one that is oriented to reference even more squarely and utterly than ordinary language. Ontological concern is not a correction of a lay thought and practice; it is foreign to the lay culture, though an outgrowth of it. (1981: 9)

When we are concerned with ontological questions, or metaphysical questions more generally, we cannot simply examine our beliefs in the terms in which we are at first prone to express them. Our beliefs must, rather, be cast into a standard notation, which will let their presuppositions shine forth. Ontology, as Quine interprets it, thus presupposes regimentation: it is only insofar as we conceive of our knowledge as cast in regimented notation that it makes sense to raise ontological questions.

The artificiality of ontology, in Quine's view, is an important point. He is often criticized for distorting ordinary thought, or doing violence to our supposed "intuitions." On his view these criticisms miss their mark entirely. More positively, he takes part of the task of ontology to be that of showing just what our theories really commit us to. Here he has an interest in ontological economy. In many cases our theories seem to commit us to accepting entities of a certain kind, but artful re-construal of the theories shows that in fact we need not presuppose entities of the given sort. Here artificiality is inevitable.

In this spirit, Quine takes the reduction of numbers to sets, and of ordered pairs to sets, to be clear philosophical achievements. (Section 53 of Quine 1960 is entitled "The Ordered Pair as Philosophical Paradigm.") In each case the reduction shows how our discourse could be rephrased so as to avoid commitment to a kind of entity to which it appears to be committed. The rephrasing is not meant as a practical substitute for normal arithmetical or set-theoretic language. Nor is it claimed that the paraphrased version gets at the "real meaning" or hidden structure of the original. It is simply that the reduction shows us how we could do everything that we need, for scientific purposes, without presupposing that there are ordered pairs (or numbers), and hence that we need not take ourselves to be committed to the existence of such things. Nor is it only towards such technical subjects as arithmetic and set-theory that Quine takes this attitude. On the contrary; he takes the same view everywhere. He holds, for example, that we can eliminate minds in terms of bodies, just as we can eliminate numbers in terms of sets. Instead of speaking of a person's mind at a given moment, we could instead speak of his or her body at that moment. Mentalistic predicates, such as "is thinking of Vienna" persist, but no exclusively mental *objects*. The maneuver is trivial, but not, in Quine's view, any the worse for that; he takes it as showing that we have no ontological commitment to minds, over and above bodies.

Quine thus stresses the artificiality of ontology, and the use of paraphrase or re-construal to eliminate apparent ontological commitments. His tactics here presuppose something implicit in our earlier discussion of knowledge, especially of observation sentences. In his view language is, of course, referential – it is *about* things. We refer to people and other objects in almost everything we say. But the relation of reference is not our fundamental cognitive relation to the world. The fundamental relation, the way that language gets to be about the world at all, is the relation of observation sentences to patterns of stimulation. (In particular, the relation of a given observation sentence to the patterns of stimulation on the basis of which speakers of the language are inclined to accept or reject the given sentence.) This is not a relation of reference, for observation sentences are not *about* patterns of stimulation. Reference is a derivative relation. Indeed the very notion of an object, that to which we refer, is in Quine's view derivative – though it comes so naturally to us that it seems inevitable. If there were a language consisting only of observation sentences, there would be no reason to attribute any reference to that language. The vastly greater complexity of our language requires inferential links between sentences, links which we can grasp only by attributing structure to the sentences. Seeing sentences as divided into terms, some of which refer, is part of this process:

Reference and ontology recede thus to the status of mere auxiliaries. True sentences, observational and theoretical, are the alpha and the omega of the scientific enterprise. They are related by structure, and objects figure as mere nodes of the structure. (Quine 1990: 31)

The derivative character of ontology manifests itself most dramatically in the Quinean doctrine known as *ontological relativity*, or the *inscrutability of reference*. Quine's claim is that it would be possible to carry out a large-scale re-construal of our knowledge, replacing every referring term by another. We could, for example, replace every term referring to a physical object by a term referring to its space-time complement, that is, to the whole of space and time *other* than the given object. Along with that we would re-construe our predicates, so that a predicate true of a given object would now be taken as true of its space-time complement. These re-construals would cancel out, leaving the truth-value of each sentence the same as before. The result would be a system of knowledge exactly like our own in its structure, including its relation to evidence, but in which each term refers to objects different from those to which its unreconstructed analogue refers. Nothing, Quine claims, prevents such a re-construal. The passage quoted in the previous paragraph continues: "What particular objects there may be is indifferent to the truth of observation sentences, indifferent to the support they lend to theoretical sentences, indifferent to the success of the theory in its predictions" (Quine 1990: 31).

Quine claims that this sort of wholesale re-construal of our knowledge is possible, that it would not be inconsistent with any part of it. True, I can and would stoutly maintain that my word "rabbit," say, refers to rabbits and not to their space-time complements. But *all* my uses of the word, including those uses in which I insist that it is rabbits that I am talking about, are open to re-construal. From the possibility of this sort of re-construal, Quine sometimes infers that there is no fact of the matter about reference – no fact of the matter as to whether my word "rabbit," say *really* refers to rabbits, or rather to the space-time complement of all the rabbits. The thesis of ontological relativity, however, is perhaps best thought of as a rejection of this idea of "really refers," insofar as it outruns the ordinary idea of reference. Taking our own language for granted we can say to what the words of another language (or, trivially, of that same language) refer; while we stay within that language we use it to refer with no more than the usual difficulties or ambiguities. Ontological relativity is, Quine says, "unproblematic but trivial" (Hahn and Schilpp 1998: 460). Only when we translate, or map our language onto itself, does the ontological relativity emerge. And here it should perhaps be seen as a reminder of the derivative status of reference.

Quine has come to give increasing emphasis to the doctrine of ontological relativity. Should we see it as conflicting with, or undermining, his insistence on realism? Since it leaves truth-values unaffected, it does not affect his claims to be a realist about truth. That latter claim, as we saw, was also one that might be doubted. Quine's claim there is perhaps that he is as much of a realist as it makes sense to be: that there simply is no coherent sense of realism stronger than his. Similarly here, perhaps. Quine claims that ontological relativity is beyond doubt; it "admits of trivial proof," he says (Hahn and Schilpp 1998: 728). So there is in his view no chance of defending a version of realism that denies it. Ontological relativity might thus be taken as showing us what realism

can come to. Ontology simply is derivative upon truth, and given any system of truths there simply will be more than one way of construing its ontology. If this undermines our previous conception of realism then, from Quine's point of view, so much the worse for that conception.

Conclusion

I shall not attempt to summarize what is already a very compressed treatment of Quine's thought. It is perhaps worth saying, however, that I have ignored, or treated very briefly, a number of issues that have occupied considerable space, both in Quine's own writings and in the works of commentators and critics. Notable examples here include the indeterminacy of translation, Quine's views of indirect discourse and of modality, the question of his physicalism, and what it amounts to, and his substantive views on ontology. My decisions as to how to use the limited space available here are of course based on my view of what is most important in Quine's thought. But the reader should perhaps know that others might have made these decisions rather differently.¹

Note

- 1 For their comments on earlier drafts of this essay I am indebted to Bill Hart, Peter Hacker, Gary Kemp, and Al Martinich.

Bibliography

Works by Quine

- 1936: "Truth by Convention," in *Philosophical Essays for A. N. Whitehead*, ed. O. H. Lee, New York: Longmans. (Reprinted in Quine, *The Ways of Paradox*, rev. edn., Cambridge, MA: Harvard University Press, 1976.)
- 1960: *Word and Object*, Cambridge, MA: MIT Press.
- 1961: *From a Logical Point of View*, 2nd, rev., edn., Cambridge, MA: Harvard University Press. (First published 1953.)
- 1963: "Carnap and Logical Truth," in Paul A. Schilpp (ed.) *The Philosophy of Rudolf Carnap*, La Salle, IL: Open Court. (Reprinted in Quine, *The Ways of Paradox*, 1976.)
- 1969: *Ontological Relativity and Other Essays*, New York: Columbia University Press.
- 1974: *Roots of Reference*, La Salle, IL: Open Court.
- 1981: *Theories and Things*, Cambridge, MA: Harvard University Press.
- 1987: *Quiddities: An Intermittently Philosophical Dictionary*, Cambridge, MA: Harvard University Press.
- 1990: *Pursuit of Truth*, Cambridge, MA: Harvard University Press.
- 1991: "Two Dogmas in Retrospect," *Canadian Journal of Philosophy* 21, pp. 265–74.
- 1995: *From Stimulus to Science*, Cambridge, MA: Harvard University Press.
- 2000: "Response to O'Grady," *The Proceedings of the World Congress of Philosophy, Boston, 1998*, Bowling Green, OH: Philosophy Documentation Center.

Works by other authors

- Carnap, Rudolf (1934) *Die logische Syntax der Sprache*; Vienna: Springer Verlag. (Published, with additions, as *The Logical Syntax of Language*, trans. Amethe Smeaton, London: Routledge and Kegan Paul, 1937.)

- Creath, Richard (ed.) (1990) *Dear Carnap, Dear Van*, Berkeley and Los Angeles: University of California Press.
- Gibson, Roger (1982) *The Philosophy of W. V. Quine*, Tampa: University Presses of South Florida.
- Gregory, Paul (1999) "Quine's Natural Epistemology," Ph.D. dissertation, University of Illinois.
- Hahn, Edwin and Schilpp, P. A. (eds.) (1998) *The Philosophy of W. V. Quine*, expanded edition, Chicago and La Salle, IL: Open Court Press. (Previously published 1986.)
- Hylton, Peter (1990–1) "Translation, Meaning, and Self-Knowledge," *Proceedings of the Aristotelian Society* 91.
- Lee, Harold N. (1998) "Discourse and Event: The Logician and Reality," in Hahn and Schilpp, pp. 295–315.
- Smart, J. J. C. (1969) "Quine's Philosophy of Science," in *Words and Objections: Essays on the Work of W. V. Quine*, ed. Donald Davidson and Jaakko Hintikka, Dordrecht: Reidel, pp. 3–13.
- Whitehead, Alfred North and Russell, Bertrand (1910–13) *Principia Mathematica*, 3 vols, Cambridge: Cambridge University Press. 2nd edn. 1927.