

PREFACE

This collection originated at a conference organized by the department of the philosophy of language and philosophy of science of the Philosophical Institute of the Czech Academy of Sciences held in Karlovy Vary in August 1995. Most of the papers in this volume were presented on that occasion. Professor Quine presented the paper included here and also commented on those given at that time. The questions for Professor Quine put forward in the present volume are arranged into three sections. They bear on issues in the theory of knowledge, the philosophy of language and the philosophy of logic.

To a some extent the organization parallels the treatment Quine himself gives to his views in some recent works. Quine's naturalizing of questions of epistemology involves a naturalistic account of the language employed in recording our knowledge claims and this in turn leads to questions about language and meaning, and logic and reference. The initial essays on naturalized epistemology deal with questions concerning the role of observational factors, which Quine addresses in his paper on Triangulation and which Szubka and George then take up. Grayling and Lehrer question the very project of naturalizing epistemology. Gibson considers Quine's relationship to Wittgenstein, Bergstrom an empiricist definition of truth, Miscevic the status of the a priori and Gjelsvik offers a naturalistic account of decision making.

The section on the philosophy of language concentrates on Quine's well known conjecture of the indeterminacy of translation. Segal, Antony, Horwich, Pagin and Stoutland examine the conjecture from varying vantage points: Segal suggests that a mentalistic semantics can mitigate against the indeterminacy and Antony that data concerning language acquisition argue that same result, Horwich focuses on whether translation provides the way to approach meaning, Pagin on the public nature of our knowledge of meanings, and Stoutland on the normative nature of rules of language.

The last part on logic and reference begins with an examination of the problem concerning vacuous names which Quine dubbed Plato's beard. Parsons and Woodruff then take up matters concerning vagueness. Neale, Ray, and Recanati raise questions concerning issues surrounding opacity: Neale and Ray on modality and Recanati on propositional attitudes.

We, and most particularly I, Alex Orenstein, would like to express indebtedness to Queens College of the City University of New York for a presidential research award which provided time to work on this project, to the Faculty Research Award Program of the City University of New York for a grant which aided us in preparing the manuscript for publication, to Wolfson, to St. Anne's and to Exeter Colleges, Oxford for the use of their facilities during much of the preparation of this work and to Roger Gibson and Douglas Quine for their help with the proofs. Last of all, we would like to record our gratitude to Marx Wartofsky, whose untimely death prevented our thanking him in person. His help made it possible for this collection to appear in the *Boston Studies in the Philosophy of Science* series.

Alex Orenstein and Peter Kotatko

I, YOU, AND IT: AN EPISTEMOLOGICAL TRIANGLE

The triangle I have in mind is one that Donald Davidson has occasionally invoked. It has you at one vertex, some one of you. I am at another vertex, and at the third vertex there is some object. The object is reflecting light at both of us, but in somewhat different patterns owing to our difference in perspective. The light rays trigger our nerve endings, setting off one train of events in your nervous system and another in mine.

The object at the third vertex is a creature, unfamiliar to me. You tell me a name for it: aardvark. What went on in your nervous system and mine when we observed the aardvark differed in perspective and probably in more. We are differently wired. Our sensations may have differed too, if that makes sense. All we clearly shared was the distal cause of our neural events, the aardvark. Still I end up associating this same word with my stimulus, my neural intake, as you did with your different intake – numerically different certainly, and somewhat different in further ways. We thus differ in the proximal causes of our concordant use of the word, but we share the distal cause, the reference, farther out on our causal chains.

This is something worth wondering about. The causal chains from the aardvark to you and to me part company already at the aardvark itself and make their divergent ways into our unlike nerve nets, but we both end up calling the creature an aardvark. We shall find that what forges our link in the triangle, linking us, is a preestablished harmony.

I have now learned the word “aardvark”, thanks to you, and I proceed to apply it to similar creatures as they turn up. You do likewise, independently, and we find ourselves in agreement in our continuing use of the term. This suggests that we are both good judges of similarity, in some objective sense of the term. But wait: I know no objective sense of the term, apart from geometry.

Subjective similarity, yes. We can make good psychological sense, experimentally testable, of perceptual similarity of neural intakes for an individual. We can test it by reinforcing and extinguishing his responses. Stimulus *A* is perceptually more similar to stimulus *B* than to *C* if *A* evokes a response that was rewarded on the occasion of *B* but penalized on the occasion of *C*.

Back now to the aardvarks. The next time one turns up and you and I are there to observe it, I say “There’s an aardvark” and you agree. My neural intakes from the two aardvarks were perceptually similar by my lights, and your two were perceptually similar by yours. There is no presumption of perceptual similarity between your intakes and mine; intersubjective perceptual similarity is not even defined.

What we have rather is a parallelism, a preestablished harmony, between your standards of perceptual similarity and mine. In general, if external events out there at the third vertex on two occasions produce neural intakes in both of us, and yours are perceptually similar for you, mine are apt to be perceptually similar for me.

Without this harmony, my perceptual similars might not have paired off so nicely with yours. My later intake might not have qualified as perceptually similar for me to my earlier one, though your later one did so qualify for you. Our scales of similarity might have diverged. Thanks to the harmony, they did not. Aardvarks look pretty much alike to me, and they look pretty much alike to you.

The preestablished harmony is needed to account for our meeting of minds not only on aardvarks, but also on what to call them: on the mellifluous Dutch disyllable "aardvark" itself. The phonetic constancy of a word, from one utterance of it to another, is itself a product of the speakers' subjective standards of perceptual similarity. Thanks to the harmony, communication proceeds apace. Oh, we sound alike. Oh, who says so? Each of us, by his own standards of perceptual similarity, all of which are in harmony.

Perceptual similarity is vital not only for language, but for all learning, all habit formation, all expectation. It was its role in learning, or habit formation, that underlay the experimental test of perceptual similarity itself. Since it is presupposed in learning, moreover, perceptual similarity cannot itself have been learned, not all of it; some had to be innate, though it gets overlaid and changed as learning progresses.

So perceptual similarity standards are in part innate and are in preestablished harmony. Natural selection accounts nicely for both traits. The relevance of similarity standards to survival becomes evident when we consider the role they play in expectation. We tend innately to expect similar events to have sequels that are in turn similar to each other. This is primitive induction. There is survival value in successful induction, successful expectation: it expedites our elusion of predators and our pursuit of prey. Natural selection, then, has favored similarity standards that mesh relatively well with the succession of natural events. This explains why our expectations have run a better than random score of success. It also explains the preestablished harmony: the standards are largely fixed in the genes of the race, the species.

What I call an observation sentence, for a given speaker, is any expression, such as "aardvark", to which he has learned to assent in response to all neural intakes within some range of perceptually similar intakes. The expression may be a single word, as here, or it may be a grammatical sentence: "It's cold", "It's raining". They are "about" the external world, not about sense data or stimulation. They are directed at the "it" vertex of the triangle, where reference lies.

An observation sentence is an *occasion* sentence, true on some occasions and false on others. An occasion sentence qualifies as an observation sentence for a speaker if he is disposed to assent to it outright when appropriately stimulated. It is an observation sentence for a society of speakers if it is an observation sentence for each member and if, further, all members who witness an occasion are disposed to agree in assenting or in not assenting.

The child's first cognitive utterances are observation sentences – thus “Mama”, “Milk”, “Dog” – though it will be a while before the child gets a grip on the schematism of space, time, and recurrent enduring objects, to which these observation sentences as we know them are committed. These adult refinements make sense of sameness of a recurrent object and distinctness of an exactly similar object. For an infant the observation sentence “Ball” would merely signal a recurrent state of affairs, at first, like “Thunder” or “It's cold”. Only later does it make sense to wonder whether it is the same ball next time around or another one like it. Only then is reification in full flower. But the reference is out in the world early and late, at the “it” vertex of the triangle. The child merely has to learn more about it. Meanwhile the range of neural intakes prompting the child's assent to the observation sentence “Ball” can remain unchanged.

At the primitive or infantile stage the observation sentences are the human analogues of the apes' cries, the signals by which an ape alerts his kin to the approach of a predator or the prospect of some fruit at a treetop. Language is, itself perhaps a luxuriant outgrowth of the apes' cries. In any event the child promptly outstrips the apes, amassing new observation sentences hand over fist and learning ways of combining them with help of prepositions and conjunctions to form further observation sentences.

He then somehow caps this achievement with what I picture as his next great acquisition, the *observation categorical*. It is a generalized expression of expectation, the fruit of induction. It is a way of joining two observation sentences to express the general expectation that whenever the one observation sentence holds, the other will be fulfilled as well. Examples: “When it snows, it's cold”; “Where there's smoke, there's fire”; “When the sun rises, the birds sing”; “When lightning, thunder”. They are our first faltering scientific laws.

The survival value of the apes' cries, and of our ordinary observation sentences, lay in vicarious observation. One learns from one's kinsman, or kinsape, about something that one couldn't see from where one was sitting at the time. Now the observation categoricals bring us much more: they bring us vicarious habituation, vicarious induction. We get the benefit of generalized expectations built up over the years by some veteran observer or even by his informant, long dead. Observation categoricals can be handed down.

There is more to be said to the glory of the observation categoricals: they are the empirical checkpoints of scientific theory. If a theory implies an observation categorical, we can subject the theory to an experimental test by bringing about fulfillment of the observation sentence that is the protasis of the categorical and then watching for fulfillment of the apodasis. I see in this the distilled essence of the experimental method, and the only empirical check to which science is susceptible. Failure of the test brands a theory as false. Any appeal must rest on challenging the observations themselves, whereof more anon.

Language and science go on developing hand in hand, in the child and in the race. I already remarked on the big ontological leap achieved in conceiving of an enduring body, a body intermittently observed between unseen trajectories. Another big ontological leap was the reification of numbers and other abstract

objects. It became clearly recognizable only with the advent of relative pronouns, or variables.

Observation sentences were at first acquired at mother's knee, and affirmed on appropriate occasions without peradventure. The growth of theory, however, calls upon us to think twice. "That's a rabbit", affirmed under appropriate stimulation, had seemed final. But at a later stage, when we allow for the identity over time of an intermittently observed body, we may come to suspect that yesterday's supposed rabbit was just this toy that we now find in the garden.

We continue to accrete observation sentences throughout life. Some, like the child's early ones, are learned outright as wholes by direct conditioning to appropriate ranges of neural intake. Mostly such acquisitions will subsequently become infected with theory in varying degrees. Other observation sentences may be assembled from words learned in earlier contexts. Speakers vary in how they arrive at the same usage of an observation sentence.

There are sentences that are learned only through theory, but become observational later in the specialist's career; thus "There was copper in it", said by the chemist after a glance at the solution, or "There goes a hyperthyroid", said by a physician after a glimpse of a stranger's face.

What I have said thus far will have been pretty familiar to those of you who have been reading my writings. But there now comes a change, the effect of correspondence with Professor Bergström. He propounded what he called an empiricist theory of truth, which I found both appealing and disconcerting. In the end I did not adopt the theory, but I came to this decision only by sharpening my own views in a quarter to which I had hitherto given too little attention: the *recanting* of observation sentences.

The speaker assents outright to "Rabbit" and retracts his assent later on discovering that what he had seen was a toy. In *Word and Object* (p. 44) I treated such intrusions of error as showing that the retracted sentence was not purely observational; theory had intruded. Theory had indeed, but I now see a better way of sorting matters out.

Immediacy of assent is still my criterion of observationality, but emphatically and unequivocally so. The assent may be recanted, but its susceptibility to recantation becomes an independent dimension, *theoreticity*, for separate consideration.

The range of neural intakes to which the speaker's assent to the sentence is keyed will of course be vague along the edges. The speaker may hesitate over "It's raining" in a fine mist, and over "That's a swan" in the startling presence of a black one. Vagueness of boundaries invests language at every turn, and I shall continue to take it in stride. But assent must be immediate when the stimulation is in the clean-cut range between vaguenesses. Immediate but fallible.

As always, the subject matter of the observation sentence is the real world, the "it" of the triangle, and it is the domain of fallible scientific theory. "It's raining", affirmed in full view of a rain-drenched window, is recanted on spotting the garden hose, and "Rabbit" on spotting the toy. Such recantation reflects conflict of observations through theoretical connections between observation sentences. The degree

of susceptibility to recantation measures how theoretic the observation sentence is. It is its *degree of theoreticity*.

An observation sentence that is perhaps minimally theoretic is "This looks blue". I say "looks" here, rather than "is", to allow for the possibility that reflected light or environmental contrast may be affecting the color that the object would otherwise show. The reference of the sentence is still to an external patch or body, or is to become so with the flowering of reification.

This move brings me to terms with Kuhn, Hanson, and Feyerabend, who have objected to observation sentences that they are suffused with theory. They are indeed, I now agree, in all degrees, but their observational role remains. The immediacy of assent shows that in one way or another they have become conditioned as wholes to ranges of perceptually similar neural intakes and so serve as observation points in the experimental testing of hypotheses. The experimenter is still free, indeed advised, to examine the circumstances of his observation for possible experimental error, or even to try again.

Sophisticated observation sentences, such as the ones about copper and the hyperthyroid, are apt to be reducible to more primitive ones, delineating more directly sensory evidence. These will tend to be less theoretic by the stated criterion, that is, less susceptible to recantation. Some sophisticated ones, however, are not thus reducible. I think of the subtle traits that the wine expert learns to detect.

Such reduction, where possible, bolsters scientific theory; for the increased resistance to recantation of observation sentences increases their dependability in checking scientific theory.

A thoroughgoing reduction project of the kind, however, is surely a forlorn hope. It is utterly alien to what goes on and went on in the development of language and science in the child and in the race. Observation sentences, already theoretic in varying degrees, are learned outright and helter-skelter by direct holophrastic conditioning. Further ones are synthesized along the way from bits of those at hand. It is not in general a matter of basing more theoretical ones on less theoretical ones. They vie with one another in a surging equilibrium of evidential claims. Such, in my sometime collaborator Joe Ullian's vivid phrase, is the web of belief.

Only limited reduction of overall theoreticity, then, is to be hoped for by reduction of some observation sentences to others. Some reduction may also be sought, however, in another way: by comparing several observation categoricals that are deducible from a theory as checkpoints. One of these may contain less theoretic observation sentences than others, and we could choose it for testing.

The last traces of analogy between my naturalized epistemology and the phenomenalist's dream of a reduction to sense data have now, it seems, been dissipated. Observationality becomes a simple behavioral matter of prompt response to stimulation, and overall reduction of the more theoretic to the less theoretic ceases to figure as an ideal.

My observation categoricals were always a schematic caricature of what goes on in testing scientific theory. They represented an adverse pair of observations as refuting a theory beyond redemption, whereas real life is less clean-cut. Now my recognition of theoreticity as a second dimension for observation sentences

QUINE AND DAVIDSON ON PERCEPTUAL KNOWLEDGE

One of the main differences between Quine's and Davidson's theories of knowledge and mind lies in their accounts of the content of perception and the way in which it contributes to our knowledge of the external world. Both thinkers are very sensitive to these differences and it has been the subject of discussion between them in recent publications. To put it very roughly, Quine holds firmly to the position that although we finally manage to get veridical knowledge of the external world, the content of our perceptions are just the triggerings of our sense receptors that give us reliable clues about the objects and happenings in our environment. Davidson considers this view to be a naturalized successor of an older defective empiricism which should be abandoned. In its place he proposes an externalist theory of perceptual content, according to which content is fully determined or constituted by the objects and events in the external world. This move, among other things, bypasses many of the troubles that Quine's approach faces and gives a solid ground for our intersubjective communication. In other words, if the central concern of Quine's epistemological project is epistemology naturalized, so the central concern of Davidson's corresponding project is epistemology externalised.¹

In the present essay I shall outline, in the first part, Quine's view of perceptual knowledge, and subsequently, in the second part, Davidson's arguments against it and his own positive account of perceptual content. In the third and final part I shall try to say why I find both of these options unacceptable, and then make some suggestions towards what I take to be a more satisfactory solution.

1. NEURAL INTAKE AND OBSERVATION SENTENCES

Quine believes that someone who has given up the idea of philosophy as independent of and prior to science, and who accepts its continuity with science, has only one available option in epistemology: she must endorse a non-mentalistic version of empiricism; one transformed into the physics of stimulus and response. Of course, this does not mean that she is restricted, while doing epistemology, to rely exclusively on the theories and hypotheses of contemporary physics; by no means – she can draw upon such disciplines as neurology, psycholinguistics and evolutionary genetics. That is, the task of a scientific or naturalized epistemologist is to adhere to the ancient slogan *nihil in mente quod non prius in sensu* but to give it a distinctive and broadly understood physicalist turn.

The assumptions from which the naturalist epistemologist starts would pass today as hardly disputable truisms or platitudes. We are highly complicated living

organisms, animals of some sort, interacting with the environment in various ways. How do we get information about this environment and the external world? Quine writes:

Science itself teaches that there is no clairvoyance; that the only information that can reach our sensory surfaces from external objects must be limited to two-dimensional optical projections and various impacts of air waves on the eardrums and some gaseous reactions in the nasal passages and a few kindred odds and ends.²

And more recently:

Our avenue of continuing information about the world is the bombardment of our sensory surfaces by rays and particles, plus some negligible kinaesthetic clues to the ups and downs of our footpath.³

The result of these causal impacts or bombardments by rays and particles are irritations of our surface or, more precisely, the triggerings of our exteroceptors. Confining our attention to a single subject triggered on a given occasion we can define the subject's global stimulation, or neural intake, as the temporary ordered set of all firing of her exteroceptors on that occasion.⁴ These neural stimulations or intakes prompt the subject to assent to or dissent from various observation sentences, sentences such as "It is raining", "That is a rabbit", or simply "Mama", "Milk", "Cold" etc. These sentences play a central role both in learning a language and in testing our scientific theories about the world.⁵

When we turn to the issue of observation sentences and assent to them, we seem to leave the safe domain of commonplace scientific truths and enter a rather murky area, the field of real philosophical difficulties. One of the problems here is to get an account of observation sentences which ensures that they successfully meet the following requirements: (i) each such sentence should be associated affirmatively with some range of one's stimulations and negatively with some other range; (ii) each such sentence should command the subject's assent or dissent outright, on the occasion of an appropriate stimulation, without further investigation and independently of her interests, etc. (iii) each observation sentence *must* command the same verdict from all linguistically competent witnesses of the occasion.⁶ The first two conditions appear to be not very difficult to meet, provided that we keep in mind the adjective "appropriate" in the phrase "appropriate stimulation" (there can, of course, be cases when the subject confronted with a black swan cannot, without further investigation, say "That is a swan" or "That is not a swan") and are able to distinguish from the wide class of observation sentences for various groups of speakers, with various scientific backgrounds, etc. a subclass of very simple and unsophisticated observation sentences common for all speakers of a given linguistic community. But the third condition, i.e. the condition of intersubjectivity, is not very easy to meet. This is because Quine has to get intersubjectivity out of the subjective and idiosyncratic stimulations of each particular member of the community. The quick reply would be to say that these stimulations are not subjective in the Cartesian sense: they are successfully investigated by the intersubjective methods

of neuroscience. But this reply is scarcely satisfactory if the observation sentences are supposed to be “the entering wedge in the learning of language”. The very idea of a mother making recourse to complicated neurological tests in order to settle whether her child associates with a given observation sentence a stimulation of the right kind, strikes us as totally absurd. So the resolution to this problem certainly does not rely on taking the subjectivity of individual stimulations as merely apparent.

Quine has become increasingly aware of this difficulty. In *The Roots of Reference* he tries to remove it by the following account of observation sentences:

A sentence is observational insofar as its truth value, on any occasion, would be agreed to by just about any member of the speech community witnessing the occasion. This definition depends ... on the idea of membership in the speech community, but that presents no problem; we can recognize membership in the speech community by mere fluency of dialogue, something we can witness even without knowing the language.⁷

The account, as it stands, suffers from the ambiguity of the word “occasion”. But further clarifications and the context of his whole epistemological project clearly suggest what Quine has in mind here: the occasions are simply sensory impingements, surface irritations, neural intakes, etc. Another defect of this account is more troublesome, namely the essential dependence of the definition of an observation sentence on the idea of membership in a speech community. Can fluency of dialogue function as a criterion or even a reliable indicator of belonging to the same speech community? Hardly. For example, Poles visiting the Czech Republic often engage in fairly fluent conversations (about matters of daily life) with the people living there, but this does not prove that they belong to the same speech community. There is only some affinity between their languages. So further refinements of the fluency criterion are needed. However, doing this membership in the same speech community comes to be based on agreement about which observation sentences are assented to or dissented from on any given occasion. If so, the above account becomes uninformative and circular.⁸

In his more recent writings (from the eighties and nineties) Quine explicitly recognizes that this way of defining observation sentences is faulty beyond repair. There is no other choice but simply to be satisfied with an account that says which sentences are observational for each individual speaker. Here is the passage applying this strategy:

An observation sentence is an occasion sentence that the speaker will consistently assent to when his sensory receptors are stimulated in certain ways, and consistently dissent from when they are stimulated in certain other ways. If querying the sentence elicits assent from the given speaker on one occasion, it will elicit assent likewise on any other occasion when the same total set of receptors is triggered; and similarly for dissent.⁹

When, however, a given sentence is observational not just for a single speaker but for a whole community (that is, for each and all its members)? The answer, at first sight, is simple: when it is observational for each individual member of the commu-

nity. And maybe this answer is *really* simple and easy to apply in deciding which sentence is observational, and which not, for someone who has extraordinary cognitive powers and capacities that enable her to look in the brains of all the speakers concerned and see whether they assent to a given sentence only when their sensory receptors are stimulated in a certain determinate way. But for ordinary speakers with ordinary cognitive powers and capacities this answer seems to be rather useless.

But Quine would reply that ordinary speakers are not in such an impoverished situation. They have empathy, a “gift of human nature”, that enable them to correlate, however fallibly, their observation sentences with the observation sentences of other members of the speech community. Empathy plays, for instance, an indispensable role in learning a language. The mother assesses the appropriateness of her child’s observation sentences by taking into account the child’s orientation and how the environment looks from there. And this use of empathy extends to various sorts of interaction between humans. To put it in Quine’s words: “We all have an uncanny knack for empathizing another’s perceptual situation, however ignorant of the physiological or optical mechanism of his perception”.¹⁰

So far I have been considering only one problem of Quine’s naturalized epistemology: the problem of defining observation sentences without abandoning objectivity. The other, which I merely mention now and to which I shall return briefly in the third part, concerns the nature of the relation of the triggerings of sense receptors or neural intakes to observation sentences. Is this a brute non-rational causal connection that makes us just happen to assent to various observation sentences? Or, is this not just a causal but *also* an epistemic relation where neural intake is evidence for or against a given observation sentence? Quine’s writings seems to be rather vague on that matter.

I turn now to arguments Donald Davidson has put forward against Quine’s views of content, perception and observation sentences, and his own positive account of these epistemological issues.

2. EXTERNAL CONTENT AND OBJECTIVE MIND

Davidson presents two arguments against Quine’s epistemological views: (i) from the nature of justification, and (ii) from the threat of scepticism. Here is the first argument.

Let us suppose, says Davidson, that we have sensations, in the form of neural intake or stimulation, which justify our observation sentences, or more generally, our beliefs about the external world. According to that picture, having the sensation of seeing a green light flashing, or having the neural intake which amounts to that, may justify the belief or observation sentence that a green light is flashing. But now, Davidson asks, does the sensation or neural intake justify the belief? And he gives the following answer:

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?¹¹

The point of this passage can be put as follows: sensations or neural intakes as such do not play any role in the justification of our beliefs or sentences; only beliefs or sentences can justify or support other beliefs or sentences. We can refer to sensations or neural intakes while giving a descriptive and causal account of how people acquire their knowledge, but an individual thinker cannot put them at the bottom of justificatory relationships holding among her beliefs and sentences. Sensations or stimulations, claims Davidson, "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".¹² There is no way of transmuting a cause into a reason, and even positing various epistemological deliverances of the senses (e.g. sensations, sense data, the given, neural intakes) won't accomplish this feat. If these deliverances are to stand in logical or justificatory relations to one's beliefs or sentences, they must already be beliefs or sentences. If they are not full-fledged beliefs or sentences, they cannot serve as justificatory reasons for other beliefs or sentences. There are, of course, causal intermediaries between our beliefs or sentences and the external world, but there are no epistemic intermediaries.

Moreover, and this is the second argument, Quine's naturalized epistemology with its reliance on neural intakes leads to scepticism in much the same way as traditional epistemologies invoking such mental entities as sense data, impressions, and the like, did. This is because it shares with them the idea that empirical knowledge requires an epistemological bridge between the external world and our beliefs, sentences or theories about it. Davidson shows how exceptionally easy it is to generate for Quine's conception the old problem of scepticism concerning the senses:

[L]et us imagine someone who, when a warthog trots by, has just the patterns of stimulation I have when there's a rabbit in view. Let us suppose the one-word sentence the warthog inspires him to assent to is "Gavagai!" Going by stimulus meaning, I translate his "Gavagai!" by my "Lo, a rabbit" though I see only a warthog and no rabbit when he says and believes ... that there is a rabbit. The supposition that leads to this conclusion is not absurd; simply a rearranged sensorium. Mere astigmatism will yield examples, deafness others; little green men and women from Mars who locate objects by sonar, like bats, present a more extreme case, and brains in vats controlled by mad scientists can provide any world you or they please.¹³

So scepticism arises because we have to start with the sensory stimulations of our receptors, that also give a meaning to our observation sentences, and then infer out of this how the external world is like or what is really the case out there. And since there are no rigid connections between what happens in the world and our neural intakes (i.e. it is not the case that the same kind of neural intake can be caused only by one particular kind of events in the world), these inferences may result in false claims about the external world.

Given the function which sensory stimulations or neural intakes have in Quine's account of observation sentences, Davidson ascribes to him (with some

QUINE AND OBSERVATION*

[T]here is more to the equating of stimulations than meets the eye, or indeed perhaps rather less than seems to do so.

—W.V. Quine, *The Roots of Reference*

Observation is central to empiricism, whose leading idea is that all knowledge is ultimately conveyed to us through our perceptual experiences. Analytic philosophy is in part distinguished by its transformation of traditional philosophical questions into ones about language, and by its emphasis on the primacy of the sentence as linguistic vehicle. So it is no wonder that the notion of an *observation sentence* is central to W.V. Quine's philosophy, which in many ways represents the zenith of empiricism as pursued in the analytic tradition. A proper understanding and assessment of his philosophy depends on a satisfactory construal of that notion, one that answers to all the roles he calls upon it to play. And yet, Quine has found "the notion of observation ... awkward to analyze" and has returned to its characterization again and again.¹

In the following, I shall sketch the history of changes rung upon this notion, the internal pressures that have led to them, and the problems they raise. I shall also offer a reconstruction of Quine's most recent conception of observation sentence and assess its consequences for the joint tenability of his views on the objectivity of science and on the indeterminacy of translation. In doing so, I hope to illuminate not only Quine's doctrines, but also something about the movement of which his philosophy is in many respects the culmination.

1. ACCOUNTS OF THE OBSERVATION SENTENCE

Quine's early metaphor that likened language, or scientific theory, to a man-made fabric that touches experience only along its edges soon came to be spelled out more fully.² By 1960, in *Word and Object*, the periphery gave way to observation sentences, and the tribunal of experience to excitation of our sensory receptors, to "the triggering of ... nerve endings."³ Observation sentences were there defined as those (1) assent to which and dissent from which are closely tied to current stimulation and (2) whose stimulus meanings across speakers are equivalent. (The stimulus meaning of a sentence for a speaker is the ordered pair of stimulations of the speaker's sensory receptors that, upon being queried with the sentence, would respectively prompt her assent or dissent.)

Intuitively, the identification of stimulation with the triggering of one's receptors was meant to secure the shareability of experience, for it was assumed that the surface neural structures of different individuals are more or less homologous. And, while (1) was intended to guarantee that a speaker's response to an observation sentence is closely tied to his experience, (2) was designed to ensure that that response be the same for all speakers.⁴

Both these features answer to roles that observation sentences play in Quine's over-all view. First, observation sentences are the points at which one's theory of the world makes contact with one's experience; without such contact, empiricism is untenable. Secondly, observation sentences are the *lingua franca* of theories; without them, resolution of intertheoretic disagreement about the nature of the world would be impossible.⁵

Because Quine conceives of language learning as a kind of theory testing,⁶ these two epistemological roles have language-acquisition counterparts. First, observation sentences are those learned without the need for "scholarship,"⁷ that is without the benefit of any prior language; if some sentences could not be learned this way, he holds, language learning would be impossible. Secondly, observation sentences, or the experiences which they tag, are the *lingua franca* of learning; without them, language acquisition as Quine conceives it, in which learner and teacher exploit their common experience, would founder.

In the light of these roles of the observation sentence, its analysis in terms of stimulus meaning can be seen to express Quine's view that the ultimate conduit for justification and meaning is experience: "The 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."⁸

On Quine's view, while translation-relevant stimulation, be it by a tack or a tachistoscope, must be shareable, we are unlikely, because of what he calls "interference from within," to find even an approximate homology between the inner physiological states of speakers who are intuitively having similar experiences.⁹ But stimulation must also be such that it should not be describable without reference to the relevant speakers: a rabbit running in the distance is not a candidate for our shared stimulation.¹⁰ Experience, insofar as our detection of it in others constrains translation, should be understood neither as an internal event, however physicalistically acceptable, nor as an event involving objects beyond the speaker. Quine split the difference, as it were, and located experience midway between the two, at the surface of the observing subject's body.

Over the years, no doubt because of the centrality of the notion of observation sentence for his views, Quine has tried to address problems with his account of experience. The first expression of worry came just a few years after publication of *Word and Object*: in 1965, Quine fretted about whether his construal of stimulation as patterns of sensory triggerings had, ironically, the effect of making stimulation unshareably personal: "If we construe stimulation patterns my way, we cannot equate them without supposing homology of receptors; and this is absurd, not only because full homology is implausible, but because it surely ought not to matter."¹¹

In fact, Quine's above account of experience entails that there are no observation sentences, for no two speakers could associate equivalent stimulus meanings with a sentence. And even if this worry could somehow be allayed with regard to humans, it can be made vivid again for extraterrestrials, what with their alien sensory apparatus. It would be surprising were Quine to deny the possibility of linguistic communication between humans and any being outfitted with differently structured surface receptors, and this reveals that something has already misfired in his account about us.¹²

This matter can be put more generally. A translation, viewed abstractly, is a preserving map from one set to a second. We have not given this notion any substance, however, until we have specified what precisely is to be preserved. Without such a specification, no constraints have been placed on an adequate translation; indeed, we might say that no substantive construal of the notion of translation has yet been given. What Quine has noted is that his previous account of translation, as a map from one language to another that preserves whatever stimulus meaning can be assigned to sentences, fails this test, for equivalent stimulus meanings cannot be attached to sentences by different speakers, or even by the same speaker at sufficiently different times.^{13,14}

Quine's response to this unsatisfactory situation, in *Theories and Things* (1981), was to define observation sentence for an individual speaker and then count as observational *tout court* any sentence that was observational for all speakers of the language.¹⁵ Unlike the account presented in *Word and Object*, this does not have the undesired consequence that there are no observation sentences. It still fails, however, because it counts as observational any sentence which, though observational for each speaker, is such that speakers are disposed to assent to it in intuitively very different stimulatory conditions.¹⁶ On this conception, the important roles of the observation sentence that were intended to be guaranteed by (2) of the earlier account cannot be secured.

Quine's skepticism about the "homology of receptors" across individuals is not in conflict with his analysis of what he calls "perceptual similarity," for this is a relation that holds between an individual's "global stimuli – ordered sets of [triggered] receptors."¹⁷ Nor is it in conflict with his claim that there exists (perhaps for reasons of natural selection) a harmony between the perceptual similarity spaces of different individuals. In *From Stimulus to Science* (1995), Quine says something that might be read as an attempt to use these notions to provide a constraint on translation of observation sentences: "Within the individual the observation sentence is keyed to a range of perceptually fairly similar global stimuli ... It is thanks to the preestablished harmony, again, that they qualify as observation sentences across the community."¹⁸ The proposal might seem to be this: *S* (a sentence of Speaker₁'s language) can legitimately be translated as *T* (a sentence of Speaker₂'s language) only if (i) *S* is conditioned to a range of global stimuli that are perceptually similar for Speaker₁, and (ii) *T* is conditioned to a range of global stimuli that are perceptually similar for Speaker₂, and (iii) the perceptual similarity spaces of Speaker₁ and of Speaker₂ are in harmony with one another. *S* and *T* are linked by virtue of the fact that they are respectively associated with

ranges of global stimuli that are respectively located in two perceptual spaces that are in harmony.

This would not do, however. To say that Speaker₁'s and Speaker₂'s perceptual similarity spaces are in harmony is just to say that "If two scenes trigger perceptually similar global stimuli in one witness, they are apt to do likewise in another."¹⁹ Therefore, from (i)–(iii) we can infer that if two scenes trigger in Speaker₁ perceptually similar global stimuli to which *S* is conditioned, then they will also trigger perceptually similar global stimuli for Speaker₂; and that if two scenes trigger in Speaker₂ perceptually similar global stimuli to which *T* is conditioned, then they will also trigger perceptually similar global stimuli for Speaker₁. But all this is consistent with *S*'s being conditioned to perceptually similar global stimuli for Speaker₁ that are caused by scenes that cause in Speaker₂ perceptually similar global stimuli to which *T* is *not* conditioned. Under these circumstances, it would be intuitively inappropriate to translate the one by the other. (The problem here is reminiscent of the one that plagued Quine's account in *Theories and Things*.)

One might suggest adding an extra clause to (i)–(iii), namely (iv): the scenes that cause the range of global stimuli to which *S* is conditioned also cause the range of global stimuli to which *T* is conditioned. What this amounts to is just that *S* and *T* are to be translated only if the scenes that would lead Speaker₁ to assent to (dissent from) the first sentence would also lead Speaker₂ to assent to (dissent from) the second.²⁰

Quine, as I understand him, should not find this proposal acceptable as an analysis of what translation preserves. We saw earlier that he will not accept external causes or events as candidates for stimulation. Consequently, the present account of translation is not one that depends on anything that Quine would recognize as sensory experience. And for just this reason, the account fails to safeguard the "cardinal tenets of empiricism" that Quine holds "unassailable": for they call for an analysis of how words acquire their communal meanings, of what translation is to be faithful to, that crucially employs a notion of stimulation, or "sensory evidence." Although empiricism demands that experience be substantively implicated in the acquisition of meaning, this account of translation gives it no significant role to play.

By contrast, Quine's *Pursuit of Truth* (1992) can be read as articulating yet another approach, which is simply to grant the privacy of stimulus meanings, and to count as observation sentences those which are intimately tied to perceptual experience and which each individual would agree in assenting to, or dissenting from, "on witnessing the occasion of utterance."²¹ Instead of talking of congruence of assent/dissent behavior in the face of like stimulation, Quine there speaks of such congruence in the face of like perceptions of the relevant occasion. What Quine might mean by this will be taken up in section 4 below.

2. INDETERMINACY OF TRANSLATION

Before considering whether translation of observation sentences is indeterminate, we should say a few words about Quine's thesis of the indeterminacy of translation.

The argument for indeterminacy that pervades most of Quine's work proceeds as follows. There is a range of acceptable evidence that the translator may appeal to in choosing a manual of translation; this constitutes "the ultimate data for the identification of meanings."²² This evidence is restricted to that which Quine believes one has, in principle, at one's disposal when learning a first language, that which one has to go on when "rating [another] as a master of the language."²³ All such evidence, Quine insists, is thoroughly public. The plight of the radical translator is meant to render vivid these everyday circumstances and the publicity of the relevant evidence: "All the objective data he has to go on are the forces that he sees impinging on the native's surfaces and the observable behavior, vocal and otherwise, of the native."²⁴ Furthermore, this publicly observable evidence cannot be brought to bear on the translation of sentences taken one at a time, but instead is relevant to accepting an entire package of translation. Finally, Quine argues that the totality of "objective data" is consistent with mutually incompatible manuals of translation. He concludes from this that the choice of manual of translation is indeterminate, that is, that there is no correct choice of manual, that there is no fact of Nature that this choice must reflect.

Some (for example, Michael Dummett) have parted company with Quine on account of his assumption that we cannot in general evaluate hypotheses about the translation of individual sentences independently of hypotheses about the translation of other sentences. Others (for example, Noam Chomsky) have dissented from Quine's conclusion because they question an important presupposition of the above inference. Evidence, they claim, is not being treated here as it is in the natural sciences. There, one standardly assumes that evidence provides one with information about how it is in some determinate corner of reality. Quine, by contrast, treats the totality of "objective data" as if it were *constitutive* of the domain of facts for which, one might have thought, it is merely evidence. Facts about meaning are exhausted by the accessible evidence, on his view, for "there is only the natives' verbal behavior for the manuals of translation to be right or wrong about."²⁵ "It is the very facts about meaning," Quine holds, "that must be construed in terms of behavior."²⁶ Quine explicitly draws the contrast he sees between the nature of evidence in semantics and in naturalistic inquiry: "Dispositions to observable behavior are all there is for semantics to be right or wrong about In the case of systems of the world, on the other hand, one is prepared to believe that reality exceeds the scope of the human apparatus in unspecifiable ways."²⁷ While physical reality is typically assumed to extend beyond the results of measurement, there is nothing to semantic reality beyond the verbal behavior of speakers. Putting it yet another way, we might say that while physical facts are not usually taken to supervene on facts about measurement, Quine holds that all truths about meaning supervene on truths about verbal behavior.

This constitutivity of evidence is what Quine means when he says that verificationism is an "attractive" doctrine that supports one route to indeterminacy.²⁸ Another, more usual, expression of this is his avowal of behaviorism: "In psychology one may or may not be a behaviorist, but in linguistics one has no choice."²⁹ In part because of his *penchant* for this latter way of expressing himself,

NATURALISTIC ASSUMPTIONS

Naturalized epistemology is epistemology based on accepting the deliverances of our best current theories about the world, and premissing them in the account we give of how we get those theories. One of its principal attractions is that it allows us to make progress with other tasks in philosophy and elsewhere, unhampered by sceptical doubts. The paralysing effect of self-conscious questions about the getting and testing of beliefs – prompted in the tradition of epistemology by an acute sense of the finitary predicament not just of each putative knower taken individually, but of the collective even as it pools the results of its members' best endeavours – is solved in naturalised epistemology by the simple expedient of avoiding those questions altogether.

But naturalised epistemology has a tendency to make one uneasy. Its attractiveness can come to seem a corner-of-the-eye affair, lasting only while one's gaze is fixed elsewhere. Under direct scrutiny it appears to have two serious defects, each individually fatal to it, but interestingly linked. These are that it is circular, and that it seems comprehensively to miss the point – or if not, to duck the demand – of traditional epistemology. Without doubt these complaints, in one or another formulation, are wearisomely familiar to naturalism's proponents, but I have yet to see a satisfactory response to them, and so take this opportunity to seek one.

In Quine's view, epistemology naturalised is epistemology treated as part of science – in effect, as an empirical psychological enquiry into how we get our beliefs about nature. In an often-quoted passage from the eponymous paper that launched the debate, Quine writes: "Epistemology, or something like it, simply falls into place as a chapter of psychology and hence of natural science. It studies a natural phenomenon, viz. a physical human subject. This human subject is accorded a certain experimentally controlled input – certain patterns of irradiation in assorted frequencies, for instance – and in the fullness of time the subject delivers as output a description of the three-dimensional external world and its history. The relation between the meagre input and the torrential output is a relation that we are prompted to study for somewhat the same reasons that have always prompted epistemology: namely, in order see how evidence relates to theory, and in what ways one's theory of nature transcends any available evidence".¹

On the face of it naturalisation would seem to mark a change of focus as against traditional concerns, bearing out the charge laid by Rorty and others that Quine has substituted a causal enquiry for the justification-seeking one that the label "epistemology" traditionally denotes. If so, of course, there could be a complaint only

about bending the name to a different purpose, a very minor injury to those still interested in justificatory questions, for what's in a name? – they could continue their task and call it “Fred”. But Quine insists that there has been no change of subject matter; naturalised epistemology is truly epistemology; what it does is to improve on the traditional variety. Like the traditional variety it tackles the question of the relation of evidence to theory; but it is epistemology grown up, explaining the relation of sensory inputs to outputs of theory on the basis of empirically checkable facts about how we learn to speak of the world. It transpires that *learning to speak* is the epistemic crux for Quine, for doing so, as he puts it, “virtually enacts the evidential relation”; and he has a psychogenetic account that sketches how.² It follows that epistemology is no longer “first philosophy” in the Cartesian sense, because there is no magisterial justificatory task for it to perform. Insofar as justification is to the point at all – and it is not – it comes tangentially and free, because it comes pragmatically; total theory works, which is all that needs to be said. (I speak here of course of justification as it would be sought by what Quine asperses as a “supra-scientific tribunal”, namely, traditional epistemology. There is an internal, naturalised sense of justification, the getting of evidence for theories by observation and hypothetico-deductive reasoning which science itself, in the form of psychology, teaches us: but by “justification” in all of what follows I shall mean the ambitious variety of traditional epistemology.)

One of Quine's main motives for supplanting traditional with naturalised epistemology is that, in his view, the former attempted to get science from sensings by a reductionist project after the various manners of Hume, Russell in one of his phases, and Carnap – who is Quine's principal target in this connection. Quine famously rejects the reductionism involved, but accepts that epistemology “contains” science in the sense that science is indeed projected from data, as the output consequent to sensation. But the containment in that direction is reciprocal to a containment in the other: science contains epistemology, because epistemology is part of psychology. The reciprocity has the effect that epistemology itself turns out to be, like the rest of science, “our own construction or projection from stimulations”.³

It is this “reciprocity” that looks squarely circular (so to speak). Quine recognises this, but argues that giving up the dream of deducing science from sense data breaks the circle, or perhaps renders it virtuous. If the enterprise of effecting a reductive translation of physical language into sense datum language cannot be carried through, then the Janus face of psychology begs no questions in looking both ways.

What is the objection to this? Well, it is that Quine has taken it that the one alternative to naturalised epistemology – that is, psychology – is epistemology which involves reductive translation. His immediate target is Carnap's “rational reconstruction”, but this can serve to represent phenomenalisms in general, since it shares with them an ambition to exhaust external-world talk in experiential talk without remainder. But why should such a strategy be thought the only alternative to naturalised epistemology? And why should it be thought that rejecting it disposes of any interest in justification?

Let us note how Quine comes to think that the reductivist project is naturalism's only alternative. He starts by stating that epistemology is concerned with the foundations of science, which includes mathematics. In just the way that study of mathematics neatly divides into "conceptual" and "doctrinal" studies – that is, investigations respectively into meaning and truth – so too does the study of natural knowledge. On the conceptual side, the task is to show how natural knowledge is based on sense experience. On the doctrinal side, the task is to justify natural knowledge in sensory terms. In Quine's view the doctrinal side fails; we have not, he says, advanced beyond Hume's despair on this head. But on the conceptual side there has been progress, made possible by the technique of contextual definition.

Carnap's "heroic efforts" to effect a rational reconstruction of natural knowledge in terms of sense experience supplemented by logic and set theory, lie on the conceptual side. If successful his efforts would identify and clarify the sensory evidence for science, even if they failed to show how its possession *justifies* science; and they would deepen our understanding of our discourse about the world. But they do not succeed, because the reductive translation required does not go through: Quine's argument needs no rehearsing here. Quine therefore writes, "If all we hope for is a reconstruction that links science to experience in explicit ways short of translation, then it would seem more sensible to settle for psychology".⁴

Let us grant the point about reductive translation. Quine next assumes that naturalised epistemology and epistemologies that turn on translation between them exhaust the options. But this is just incorrect; there are a number of non-reductive ways for the justificatory enterprise to proceed. Take just one example: a remarkable feature of Quine's claim about the foundering of the doctrinal project in Hume is his neglect of what that failure immediately prompted, namely, Kant's critical enterprise, and what that directly or indirectly prompted in its turn, namely, the possibilities it suggested to some – Strawson and the very late Wittgenstein, in not very different ways, included – for exploiting more forensic conceptions of justification than is envisaged by austere versions of empiricism taken alone.⁵ It is not necessary to dilate on these options here; we have merely to be reminded of their existence to see that we are not faced with a simple disjunctive syllogism. If naturalised epistemology is to rest to any degree on ruling out the opposition, it will have to show that all other approaches to the doctrinal task founder too.

Quine resisted the charge of circularity, recall, by saying that so long as investigation of the evidential link between sense data and theory is not regarded as addressing questions about justification, no circle is at stake. He assumed that the attempt to secure justification consists in the attempt to reduce theory to data, and so could argue that abandoning hopes of a translation amounts to abandoning the justificatory enterprise. If so, the circle breaks, or turns virtuous: where there is no attempt at justification, it begs no questions to premiss in epistemology what traditional epistemology took itself as having to justify. But the foregoing remarks show that so far we have not been shown that the justificatory enterprise is bankrupt, because we are not bound to identify it with the reductionism Quine rejects.

The next step in showing that naturalism does not escape circularity is to demonstrate something stronger, namely, that the justificatory question is anyway

unavoidable. Quine says that epistemology is concerned with the foundations of science. A concern with the foundations of science – a foundational concern – sounds like a concern with assumptions, aims, ontology, methodology and reasoning, one of the chief points in investigating all which – if not *the* chief point – is to ascertain how to tell good theories from bad. This is much more than just seeking to understand the links between evidence and theory, where such understanding falls short of providing justifications. We wish to know when the links are strong enough to bet one's money or one's life on them. We recognise that there are different kinds of evidential relations, individuated by subject matter. We wish to control what can count as evidence for a given subject matter, and to know when caution (or to call it otherwise: doubt) is appropriate, and what can legitimately prompt it. Most of our concerns – in the physician's consulting room, on the construction site, in the pharmaceutical laboratory, on the battlefield – are austere practical ones, where reliable means of forming judgements are at a premium. We therefore have to try to advance beyond merely *understanding* evidential links where doing so falls short of showing how they bestow a license to rely on certain of them, for such understanding must be the basis for yielding something more: namely, norms. A conception of "getting it right" in a given area of endeavour matters to us; and it demands tests and a way of recognising when they have been passed or failed. Hence, the justificatory enterprise is a non-negotiable part of epistemology. It is why epistemology got started.

So that there is no chance of misunderstanding what is at stake here, allow me to iterate. Quine describes the enterprise of naturalised epistemology as the endeavour (I quote) "simply to understand the link between observation and science" when he is specifically seeking to avoid the circularity charge. That makes it sound as if the task is a very modest one; the passage in full reads, "scruples against circularity have little point once we have stopped dreaming of deducing science from observations. If we are out simply to understand the link between observation and science, we are well advised to use any available information, including that provided by the very science whose link with observation we are trying to understand".⁶ Now as we see, the strongly normative character that we expect to follow from understanding evidential links alters the picture. Getting into possession of epistemic norms neither *has* to, nor *does* amount to, deducing science from observations, but neither is it *just* understanding a non-justificatory connection. So we could entirely eschew the attempt to carry out remainderless translations of talk about the external world into a reducing class of sensation sentences, but the attempt to demonstrate the nature and relative strengths of the support provided to theories by observations is a practical necessity, and for this it will not do to premiss what is to be tested in an account of what tests them. Given that, the circle on which naturalised epistemology bases itself is vicious.

At this juncture, by way of aside, one might register a concern about the non-justificatory notion of "evidence" in play here. Standardly, the notion seems intrinsically to be one that exists precisely to play the traditional epistemological role of confirming and infirming, supporting and weakening given claims by its respective presences or absences. In Quine the notion is tied to sensory stimulation, and left

otherwise – and explicitly – unexplained: he writes that in his theory “the term ‘evidence’ gets no explication and plays no role”.⁷ This, obviously, is a luxury that can be afforded only when we have given up on doctrinal studies altogether.

The foregoing remarks are based on a *pro tempore* acceptance of Quine’s claim that epistemology concerns the foundations of science. But – and here is the second of the two objections to be urged against the naturalising programme – epistemology is surely concerned with a good deal more than this. It is more even than a concern with science (that is, with the superstructure as well as the foundations of science), unless the word “science” is stretched to denote with complete generality everything we hope we know. (It had this meaning in the early modern period.) For there are plenty of things we might hope to know beyond the structure and properties of the material world, which is what the natural sciences deal with. We wish also to know about history, for example, and the motives and feelings of others, and whether there is anything of value in the world, ethically and aesthetically; and we are surely interested in the differences between ways of knowing some of these things and knowing others. And as before, we have our exigent practical concerns, which makes getting things right a vital interest, and poses the demand that traditional epistemology exists precisely as an attempt to satisfy. Even the most hardened non-cognitivist in ethics, for example, must accept that we are good readers of the intentions of our fellows, and although physical entities play their part in this – raised eyebrows, air-waves issuing from mouths – it takes an even more exotic form of reductionism than Quine repudiates in Carnap, namely a belief that in the end everything will be expressible in the language of physics, to think that our other epistemic interests, in history and folk psychology and these other spheres, can be naturalised along with the epistemology of science.

Commitment to physicalist reduction cannot be foisted on Quine, however, because although it is naturalistic, it does not exhaust naturalism. Roger Gibson reminds us that Quine infers from the fallibility of science the view that future science might tell us that there is more to evidence than sensory stimulation, and that therefore adherence to empiricism must be tentative.⁸ So naturalism is not to be defined as acceptance of the deliverances of physical theory together with what it tells us about how we get it (that is, by empirical means); it is rather to be defined as acceptance of *current* such theory, for any “current”. If the theory changes, acceptance must go with it; naturalism lies, in short, in following the fashion. This is integral to naturalistic epistemology’s break with justificatory concerns, as witness these facts: a good naturalist in the 13th century AD, or the 13th century BC, would on his naturalistic principles have been bound to go with the current theory of the day. The dramatic nature of past theory change teaches us that being naturalistic at any point in the history of science is therefore a matter quite independent of whatever justification scientists in different periods have had for their claims, whether about the fifth essence, the pulsative faculty of the heart, phlogiston, or the colour and strangeness of quarks: take your pick.

Naturally we all feel confident that even if quarks go the way of phlogiston, current science has got it much righter than ever before, and in some cases, we suspect, has got it right *simpliciter*, so it at least seems safe to be epistemologically

JUSTIFICATION, COHERENCE AND QUINE

Quine has influenced epistemology beyond his own project of naturalization. His early work influenced my own work in the coherence theory of knowledge,¹ and I should like to take this occasion to raise some questions about coherence and knowledge. Quine was impressed by the failure of the reductionist program of phenomenalism. I agree with him about the failure of that program. I also agree with him that the consequence of that insight is that it is some system that confronts the world and our sensory experience of it. Knowledge, I concluded, must result from some combination of coherence and truth. I learned that from Quine but am less content to follow him down the path of naturalization. I want to explain why I took another path from his insights, a less revisionary one.

Anthony Grayling has argued that there are different conclusions that might be drawn from the failure of phenomenalism. There were, historically considered, two objectives of phenomenalism. The first was to give an account of the meaning of statements concerning material objects in terms of the language of sense data and, in that way, to reduce the meaning of material object statements to the sound foundation of sense experience. Let us call this the *meaning basing objective*. A second objective of phenomenalism was concerned, not with meaning, but with the justification of inference of material object statements from statements of sensory experience. Let us call this the *inference basing objective*. These two objectives of phenomenalism were run together historically, but they are distinct.

I have argued that success of phenomenalism in meeting the meaning basing objective would not have met the justification basing objective. The simple explanation is that those material object statements analyzed phenomenally would affirm more about sensory experience than any one person and, indeed, than all people would experience. It would tell us about future sensory experiences as well as present ones and leave ample room for the skeptic to drive a Humean wedge in the argument. Thus, the success of phenomenalism would not have been sufficient to answer a skeptic about the material world.

A similar point is made by Quine himself who remarked, assuming the success of phenomenalistic translation,

"But the mere fact that a sentence is *couched* in terms of observation, logic, and set theory does not mean that it can be *proved* from observation sentences by logic and set theory. The most modest generalizations about observable traits will cover more cases than its utterer can have had occasion to observe."²

Phenomenalism fails, and if it were to succeed it would not suffice for obtaining the justification basing objective. It is at this point that Quine decides that we should settle for psychology to provide us with an account of the relationship between sensory stimulation and the external world. He says,

“The stimulation of his receptors is all the evidence anybody has had to go on, ultimately, in arriving at his picture of the world. Why not just see how this construction proceeds? Why not settle for psychology?”³

Against the objection that using psychology to validate our claims about the external world is circular, Quine replies,

“However, such scruples against circularity have little point once we have stopped dreaming of deducing science from observation. If we are out to understand the link between observation and science, we are well advised to use any available information, including that provided by the very science whose link with observation we are seeking to understand.”⁴

Thus, Quine recognized that the success of phenomenalism was also not necessary to obtaining the inference basing objective. He might have gone on to argue that once we combine fallibilism with the assumption that it is always some scientific system of statements rather than a single statement with which we confront our experience of the world, we can use what finds its way into that system to justify our inferences. What we need to justify the inference from sensory experience to claims that exceed what we have so far experienced is some additional premise. But as we confront experience with our description of it, we do so with a background of statements that may tell us when our descriptions are justified and when they are not. It is, after all, not only perceptual statements that comprise our background system but statements about when we should trust our senses and when we should demure.

That is what I made of Quine and went on to define justification as coherence with a background system. But Quine himself took a more radical turn toward psychologism. It is only psychology, he proposes, that we should use to obtain our understanding of the relationship between sentences about sensory stimulation and sentences about the external world. Can psychology satisfy the justification basing objective? Can psychology tell us when our inferences from sensory experience to the external world are justified? It appears not, though this depends on how strictly one takes the notion of psychology. For, it appears that psychology would offer us an account how sensory experience influences the formation of beliefs about the external world while remaining silent about whether the beliefs so formed are justified. That is how I understand the point that Anthony Grayling has raised. After all, part of psychology, abnormal psychology, might tell us how the beliefs of psychotics are formed from sensory experience without commenting on whether they are justified. In fact, it appears that psychotics reason in quite normal ways, as well if not better than the rest of us, which suggests that they may suffer sensory distortion. The psychotic may, for example, believe that he observes that President

Clinton is speaking to him personally when Clinton appears on television and is communicating a secret message that only he and Clinton can understand. Now we think that these beliefs are not justified, but that conclusion is a step beyond anything that psychology tells us. Psychology tells us how beliefs originate but about whether those beliefs are justified she remains silent as a stone.

One naturalistic reply is that science can tell us when the inference will lead to a scientific conclusion that is true and that we have no better guide than science itself to tell us when our conclusions from sense experience are true and when they are not true, when, for example, the objects we think exist really do exist and when we are deluded. That is suggested by the passage above and others as well. The point is that once we give up the idea that we can validate conclusions, that is, prove them, from sensory experience, we see that we have nothing better than science to tell us whether those conclusions are true or not. And why, after all, should we want anything better than science? There is, of course, a circle in this, but that is the consequence of giving up the theory of foundational validation and is not to be deplored. In short, science cannot only tell us how our beliefs are formed from sense experience but can also use that information to sort the true from the false, or, at least, to do the best we can.

I think that there is a question remaining, however, which concerns what we are justified in believing. Quine may avoid the question and, in past discussions, has done so. Why concern oneself with the question of what one is justified in believing? Either one has an interest in truth or one does not, and if one does and thinks that science is the best way to get to it, then one follows the path of science. And so Quine does, and admirably so. Is there a question left over? There is. We may have an interest in truth, an interest in the objectives of science, and yet raise the question of whether we are justified in accepting what we do in the interest of achieving our objectives. We may appeal to other things that we accept, the contents of a background system, to answer this question. But a question remains. Are we justified in accepting any of things we do? Obviously we are justified. But why? How does justification get into the system?

This question is not a quest for a foundation but for an explanation. A philosopher who affirms that we just are justified in what we accept and that is all there is to it leaves us with kind of epistemic surd. What I want to suggest is that there is no need to leave ourselves with the surd once we have followed Quine into the perspective of allowing us to use our background system to explain why we are justified in accepting what we do to meet our scientific objectives. We need only ask ourselves what it is that we accept that is supported by other things we accept that enables us to explain why we are justified in accepting what we do.

The answer to the question is not difficult to find. We trust ourselves in accepting what we do to obtain our objectives. Moreover, we trust ourselves in preferring the objectives we do. We also consider ourselves worthy of our trust in what we thus accept and prefer. To put the matter another way, we consider ourselves trustworthy in what we accept, though acknowledging our fallibility. We think that we can be both trustworthy and fallible. Now suppose that we are trustworthy as we think we are in accepting, among other things, that we are trustworthy. That, I suggest,

explains why we are justified in accepting what we do, including that we are trustworthy. If we are trustworthy in what we accept and in our preferred objectives, then we are justified in accepting what we do for those our objectives. That does not mean that we automatically attain our objectives, that what we accept is always true, for example, for fallibilism is our cognitive burden. What it means is that there is at least a minimal level of justification in what we accept if we are trustworthy in accepting what we do to obtain our objective. And, of course, we accept that we are trustworthy.

Is our acceptance that we are trustworthy in what we accept a kind of epistemic surd? If asked why we are justified in accepting that, must we say simply that we are justified and that is all there is to it? We can do better. We can answer that we are trustworthy in accepting that we are trustworthy. The advantage of this answer is that if it is correct, it does explain why we are justified in accepting it. If we are trustworthy in accepting what we do, including that we are trustworthy, why, then, that is why we are justified. We are justified because our acceptance of what we thus accept is a consequence of our trustworthiness. If so, then our obtaining our objectives by accepting what we do is not unexplained or just a matter of luck but is instead a result of our being trustworthy as we suppose ourselves to be.

Is our acceptance of our trustworthiness a foundation? No. It is, in fact, supported by other things we accept rather than being the foundation for them. Of course, we may appeal to our acceptance of our trustworthiness to justify those other things we accept. The metaphor I prefer is that of a keystone. The acceptance of our trustworthiness is a keystone in the system of things we accept. It is supported by the other things we accept but the vaulted arch of science collapses when it is removed. It does, in a way, support itself, but only as a result of being supported by the other stones in the arch.⁵

Once we suppose that the background system of justification contains the keystone, knowledge results from coherence with the system combined with truth. But the initial justification of the contents of system is essential or the system counts for naught. Given the initial justification of what we accept, we can proceed to accept justified conclusions even though, as Quine says, we cannot validate or prove them. Why is the justification not a validation or a proof? Because, of course, circular reasoning cannot prove or validate anything. No skeptic can be expected to accept the reasoning above concerning trustworthiness if she does not concede that we are trustworthy in what we accept or in our preference for our objectives. But, though she cannot thus be proven wrong, we may be trustworthy as we think we are nonetheless. We are then justified in accepting that the skeptic is wrong, though we cannot prove or validate this claim.

Quine has argued at the end of the symposium that immediacy is something we should trust, that what we immediately accept, admitting this to admit of degrees, is a fallible guide to the truth of the matter. I suggest that he accepts that immediacy is worthy of our trust, though we are fallible. He also accepts that we are worthy of our trust in what we accept. And that, I suggest, is the keystone of our acceptance system. Our trustworthiness results, as I have noted, from our trustworthiness in the way in which we change what we accept and change our ways of changing. Is trust-

worthiness just a matter of being correct with a certain frequency? All we can say about how often we must be right in order to be trustworthy is that we must be correct a trustworthy amount of the time. That shows that trustworthiness is not reducible to anything naturalistic which, I suspect, will lead to Quine to eschew such discourse. For all that, Quine considers what he accepts to be worthy of his trust, and he has proven brilliantly justified in accepting his trustworthiness.

University of Arizona

NOTES

- ¹ Keith Lehrer, *Theory of Knowledge* (London and Boulder: Routledge and Westview Press, 1990).
- ² Willard Van Orman Quine, "Epistemology Naturalized," in *Ontological Relativity and other essays* (New York: Columbia University Press, 1969), p. 74.
- ³ *Op. cit.*, p. 74.
- ⁴ *Ibid.*
- ⁵ Keith Lehrer, *Self-Trust: A Study of Reason, Knowledge and Autonomy* (Oxford: Clarendon Press, 1997), Chap. 1.