other volumes of collected essays by donald davidson

essays on actions and events inquiries into truth and interpretation problems of rationality (forthcoming) truth, language, and history (forthcoming)

# subjective, intersubjective, objective

donald davidson

clarendon press • oxford

### oxford

university press

cireat clarendon street. oxford ox2 6dp oxford university press is a department of the university of oxford. it furthers the university's objective of excellence in research, scholarship, and education by publishing worldwide in

oxford new york

auckland bangkok buenos aires cape town chennai dar es salaam delhi hong kong istanbul karachi kolkata kuala lumpur madrid melbourne mexico city mumbai nairobi sao paulo shanghai singapore taipei tokyo toronto

oxford is a registered trade mark of oxford university press in the uk and in certain other countries

published in the united states by oxford university press inc., new york

© in this collection donald davidson 2001

the moral rights of the authors have been asserted database right oxford university press (maker)

first published 2001 reprinted 2002

all rights reserved. no part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without the prior permission in writing of oxford university press, or as expressly permitted by law, or under terms agreed with the appropriate reprographics rights organizations. enquiries concerning reproduction outside the scope of the above should be sent to the rights department, oxford university press, at the address above

you must not circulate this book in any other binding or cover and you must impose this same condition on any acquirer

british library cataloguing in publication data data

available

library of congress cataloging in publication data data available

isbn 0-19-823752-9 isbn 0-19-823753-7 (pbk.)

3 5 7 9 10 8 6 4 2

typeset in times by cambrian typesetten, **friailey**, **surrey** printed in great britain on acid-free paper by biddies ltd., guildford & kings lynn

to marcia cavell

## contents

provenance of the essays and	ix
acknowledgements	xiii

introduction

subjective	1
12. first person authority (1984)	3
13. knowing one's own mind (1987)	15
14. the myth of the subjective (1988)	39
15. what is present to the mind? (1989)	53
16. indeterminism and antirealism (1997)	69
17. the irreducibility of the concept of the self (1998)	85
intersubjective	93
18. rational animals (1982)	95
19. the second person (1992)	107
20. the emergence of thought (1997)	123
objective	135
21. a coherence theory of truth and knowledge (1983)	137
afterthoughts (1987)	154
22. empirical content (1982)	159
23. epistemology and truth (1988)	177
24. epistemology externalized (1990)	193
25. three varieties of knowledge (1991)	205
contents list of volumes of essays by donald davidson	221
bibliographical references .	227
index	233

# provenance of the essays and acknowledgements

essay 1, 'first person authority', was read at a conference on intentionality organized by henri lauener and held in biel, switzerland, in 1983. earlier versions had been read at the university of illinois at chicago circle, stanford university, and the university of colorado. it was first published in *dialectic*^ 38 (1984), 101-11.

'knowing one's own mind', essay 2, was delivered as the presidential address at the sixtieth annual pacific division meeting of the american philosophical association in los angeles on march 28, 1986, and published in *proceedings and addresses of the american philosophical association* (1987), 441-58. it is reprinted here by permission of the american philosophical association. i am greatly indebted to akeel bilgrami and ernie lepore for criticism and advice, and to tyler burge, who generously tried to correct my understanding of his work.

essay 3, 'the myth of the subjective', was read at a conference on consciousness, language, and art in vienna in 1986, and was published in the proceedings *bewusstsein, sprache und die kunst*, edited by michael benedikt and rudolf berger, by edition s. verlag der osterreichischen staatsdruckerei, 1988.

essay 4, 'what is present to the mind?', was delivered at the second international france veber colloquium, held in bad radkersburg, austria, and gornja radgona, in what was then yugoslavia and is now slovenia. it was published in a book of essays on my work, most of them delivered at that conference. the editors of the book were johannes brandl and wolfgang gombocz, and it appeared as a special volume of *grazer philosophische studien* (vol. 36, 1989) titled *the mind of donald davidson* (amsterdam: rodopi).

#### x provenance of the essays

'indeterminism and antirealism', essay 5, was read at a conference on realism and antirealism at santa clara university in early 1992. it was published in *realism/antirealism and epistemology*, edited by c. b. kulp (lanham, maryland: rowman & littlefield, 1997).

essay 6, the irreducibility of the concept of the self, was written with the help and advice of marcia cavell for a book designed to honor the work of dieter henrich. the book, *philosophic in synthetischer absicht*, was edited by marcelo stamm (stuttgart: klettcotta, 1998).

essay 7, 'rational animals', was delivered at a conference organized by henri lauener which took place in biel, switzerland, in 1981. it was first published in *dialectics* 36 (1982), 317-27.

'the second person', which is essay 8, was given as a talk at a conference on wittgenstein in paris in 1989. it was published under the title 'jusqu'ou va le caractere public d'une langue?' in *wittgenstein et la philosophie aujourd'hui*, edited and translated into french by j. sebestik and a. soulez (paris: meridiens klincksieck, 1992). in the same year a somewhat modified english version was published in *midwest studies in philosophy*, 17, which was edited by p. french, t e. uehling, and h. wettstein (indianapolis: university of notre dame press). about four pages near the end are taken from 'the conditions of thought', a paper written for, and delivered at, a plenary session at the world congress of philosophy in brighton, 1988. it was published in *le cahier du college international de philosophie* (paris: editions osiris, 1989).

essay 9, 'the emergence of thought', was given as a talk at a seminar on emergence held at the university of frankfurt in 1993. translated into german by t. marschner, it was published with the title 'die emergenz des denkens' in *die erfindung des universums? neue oberlegungen zur philosophischen kosmologie*, edited by w. g. saltzer, p. eisenhardt, d. kurth, and r. e. zimmermann (frankfurt am main: insel verlag, 1997). it was subsequently published in english with its present title in *erkenntnis*, 51 (1999), 7-17.

essay 10, 'a coherence theory of truth and knowledge', was my contribution to a colloquium organized by richard rorty as part of the 1981 stuttgart hegel congress. w. v. quine and hilary putnam were the other participants in the colloquium. our papers were published two years later in the proceedings of that congress, *kant oder hegel*? after stuttgart the four of us had a more leisurely exchange on the same topics at the university of heidelberg. when the pacific division of the american philosophical association met in march of 1983, rorty read a paper titled 'pragmatism, davidson, and truth'. in it he commented on some of the things i had written in 'a coherence theory of truth and knowledge'; his paper was subsequently published (with revisions) in *truth and interpretation: perspectives on the philosophy of donald davidson.* i replied with 'afterthoughts, 1987', which was first published in conjunction with a reprinting of 'a coherence theory of truth and knowledge' in *reading rorty*, edited by alan malachowski (cambridge: blackwell, 1990). 'afterthoughts' as printed here is, aside from some bibliographical details noted above, a reprint of 'afterthoughts, 1987'.

'empirical content', essay 11, was read at a conference in vienna (in the house wittgenstein helped design) to celebrate the hundredth anniversary of the births of moritz schlick and otto neurath. it was published in the year of the conference in *grazer philosophische studien*, 16-17(1982).

essay 12, 'epistemology and truth', was read at the fourth panamerican philosophy conference, held in cordoba, argentina, in september of 1987. it was published in the proceedings of that conference by the national university of cordoba in 1988.

essay 13, 'epistemology externalized', was read in an early version at a sadaf meeting in buenos aires in 1989. the next year it was presented at a conference in biel, switzerland, and published in 1991 in *dialectica*, 45 2-3: 191-202.

'three varieties of knowledge', essay 14, was written to be delivered in february of 1991 as an a. j. ayer memorial lecture, sponsored by the royal institute of philosophy. when the date for the lecture arrived, a snowstorm disabled the transportation system of southern england to such an extent that the lecture was cancelled. it was subsequently published with other memorial lectures in *a. j. ayer memorial exxayx: royal institute of philosophy supplement*, 30, edited by a. phillips griffiths (cambridge university press, 1991). with some changes it was delivered as a heisenberg lecture in one of the upgraded stables of the nymphenburg palace in munich and published in german in *merkur* with the same title as the title of the present book.

### introduction

the essays in this book are concerned with three sorts of propositional knowledge and the relations among them. we all have knowledge of our own minds, knowledge of the contents of other minds, and knowledge of the shared environment. the subsections of the book are titled subjective, intersubjective, and objective, the words track real differences. first person knowledge is distinguished by the fact that we can legitimately claim a unique sort of authority with respect to what we believe, want, intend, and some other attitudes. second person knowledge and knowledge of the rest of the world of nature do not have this authority, but they differ from each other in that our knowledge of other minds is normative in a way the latter is not. all three varieties of knowledge are, however, objective in the sense that their truth is independent of their being believed to be true. this is obvious in the second two cases, but it holds even in the case of beliefs about our own beliefs and other attitudes: such beliefs can be wrong, all our knowledge is also objective in the sense that it could for the most part be expressed by concepts which have a place in a publically shared scheme of things.

*essay 1,* 'first person authority', asks what explains the presumption that a speaker is right when he sincerely attributes a belief, desire, or intention to his present self, while no such presumption is appropriate when others make similar attributions to him. it is argued that 'solutions' to the problem of other minds which merely restate the asymmetry leave the field open to the sceptic. a new explanation of first person authority is offered which traces the source of the authority to a necessary feature of the interpretation of speech.

essav 2, 'knowing one's own mind', takes up an apparent difficulty about first person authority: how can we reconcile the fact that

#### xiv introduction

the contents of our minds are in part determined by external factors of which we are ignorant with the claim that we know the contents of our minds without (normally) needing or appealing to evidence? i argue that the answer depends, among other things, on giving up the idea of 'objects before the mind' none of whose attributes can be hidden from the agent.

in *essay 3*, 'the myth of the subjective', i try to make clear what it means to deny that there are objects before the mind when we have sensations or think. the idea that there is a fundamental distinction to be drawn between uninterpreted experience and an organizing structure of concepts is closely related to the supposed dichotomy of the subjective and the objective. these dualisms have dominated and defined the problems of much of modern philosophy, the problems not only of epistemology, but also of philosophy of language and philosophy of mind. in this essay i inveigh against the picture of the mind that is presupposed by the dualisms, and so against the dualisms themselves, and the epistemological and metaphysical positions based on them. the view that the subjective is the foundation of objective empirical knowledge is attacked; it is claimed that empirical knowledge has no epistemological foundation, and needs none.

*essay* 4, 'what is present to the mind?', continues the theme of the last two essays with fresh examples and arguments restated.

a number of philosophers have questioned whether there is any 'fact of the matter' concerning the propositional attitudes. *essay 5*, 'indeterminism and antirealism', attempts to put these doubts to rest. in particular it disputes the claim that if one accepts quine's indeterminacy thesis, as i do, then one has abandoned first person authority.

essay 6, the irreducibility of the concept of the self, emphasizes the features of our beliefs about our present attitudes which remain in place after we give up the myth of the subjective and its mental objects. these features include, of course, the special authority that attends such beliefs and the irreducible role of indexical sentences. it is the thoughts such sentences express which relate us and our speech to the world around us. i also discuss briefly the fact that there is no final court of appeal beyond our own standards of rationality, a point that is raised again in essay 14.

*esxav 7*, 'rational animals', is one of a number of attempts i have made to specify some of the ingredients of rationality. by rationality i mean whatever involves propositional thought. in an earlier essay ('thought and talk', essay 11 in *inquiries into truth and interpretation*) i argued for a mutual dependence of thought and language. many readers were not persuaded. here i try again, taking a different tack, one 1 have subsequently developed in the essays that follow. the considerations i adduce for the close connection between language and thought do not constitute anything like a demonstration or proof. they depend in part on what i think we know about creatures like us.

in essav 8, 'the second person', i dwell on the idea that language is necessarily social. it is argued that to have thoughts, and so to mean anything in speaking, it is necessary to understand, and be understood by, a second person. if wittgenstein held that language is necessarily social, then the central thesis of this essay is wittgensteinian. but it is denied that communication requires that one person speak as others do. rather, the objectivity which thought and language demand depends on the mutual and simultaneous responses of two or more creatures to common distal stimuli and to one another's responses, this three-way relation among two speakers and a common world i call 'triangulation'. in the end, the idea is as simple as that of ostensive learning, but with an insistence that triangulation is not a matter of one person grasping a meaning already there, but a performance that (when fully fleshed out) bestows a content on language. this thesis, and its ramifications for philosophy of mind, language, and epistemology, turn up again and again in my work after 1982 (essay 7). carol rovane, akeel bilgrami, and marcia cavell were early critics of this idea. their suggestions and candid doubts helped greatly in shaping my thinking.

essay 9, 'the emergence of thought', asks how we are to describe the transition from the prelinguistic, preconceptual mind of an infant to a child with language, beliefs, and the other propositional attitudes. i argue that we do not have a vocabulary for describing the early stages of such development, and that it is unclear what would satisfy the felt need for such a vocabulary. it is suggested that a few major steps can be distinguished by considering the strengths of the semantic theories required for various languages.

*essay 10,* 'a coherence theory of truth and knowledge', was written in 1981, before any of the first nine of these essays. there is no paper i have written i would like more to rewrite. it has understandably attracted much criticism, which is why i reprint it here

#### xvi introduction

without change. anyway, i have in effect been rewriting it ever since it was written; the nine essays that precede it in this book are partial evidence of my subsequent misgivings, and essay 14 is further testimony. i have also tried to make amends in replies to various critics in the philosophy of donald davidson (in the library of living philosophers, edited by lewis e. hahn, (chicago: open court, 1999)) and in 'replies' to barry stroud, john mcdowell, and richard rorty in critica 30 (1998). what i would most like to correct is the impression that i think experience and perception play no role in our beliefs about the world. 'experience' and 'perception' are perfectly good words for whatever it is that goes on in our minds when we look around us, smell, touch, hear, and taste. i was so eager to get across the idea (for which i should have given credit to wilfrid sellars) that epistemic intermediaries between the world and our beliefs are a mistake that i made it sound to many readers as though i were repudiating all serious commerce between world and mind. in truth my thesis then as now is that the connection is causal and, in the case of perception, direct. to perceive that it is snowing is, under appropriate circumstances, to be caused (in the right way) by one's senses to believe that it is snowing by the actually falling snow. sensations no doubt play their role, but that role is not that of providing evidence for the belief.

essay 11, 'empirical content', provides a historical background for, and commentary on, the theme of essay 10. although neurath and schlick were more enmeshed in a very old debate than they seemed to realize, their sense that they were saving philosophy from its metaphysical past gave zest and a linguistic turn to a well-worn problem.

'epistemology and truth', *essay 12*, discusses the relation of epistemology to truth. two positions are often seen as opposed: that truth is 'radically non-epistemic' (in putnam's words) or that it is to be understood in terms of what it is possible (in practice, in theory, or ideally) for us to know. neither of these alternatives, it is argued, is acceptable. truth cannot be limited to what we can or could determine to be true; nevertheless, there are firm reasons to connect truth with true belief in one way or another. a way of partially reconciling the two positions, based on the ideas explored in essay 10, is outlined.

from the time of descartes most epistemology has been based on first person knowledge. according to the usual story we must begin with what is most certain: knowledge of our own sensations and thoughts. in one way or another we progress, if we can, to knowledge of an objective external world. there is then the final, tenuous, step to knowledge of other minds. in essay 13, 'epistemology externali/,ed', i argue for a total revision of this picture. all prepositional thought, whether positive or skeptical, whether of the inner or of the outer, requires possession of the concept of objective truth, and this concept is accessible only to those creatures in communication with others. third person knowledge-knowledge of other mindsis thus essential to all other knowledge. but third person knowledge is impossible without knowledge of a shared world of objects in a shared time and space. thus the acquisition of knowledge is not based on a progression from the subjective to the objective; it emerges holistically, and is interpersonal from the start. several forms of externalism are examined and found wanting. it is argued that triangulation, which has featured in many other essays in this book, corrects and augments both a version of perceptual externalism and a version of social externalism.

'three varieties of knowledge', *essay 14*, comes closest to pulling together the main ideas in this book. if all the essays had been written after my thoughts had gelled, three varieties of knowledge' would certainly have come first, and a reader who wants an overview might well begin here.

i have tinkered with the essays in this book to improve the grammar or style and occasionally to delete repetition or what i now see as a minor mistake. i am sure many more errors persist, and there is probably too much hammering away at certain theses. but when i contemplated major rewriting i realized that either i must reproduce my past work pretty much as it stood or simply start over again, and this would take years. i console myself with the thought that a fresh start would mean wiping out what amounts to a history of my attempts to come to grips with the triangle composed of a person, his society, and the shared environment. it would also deprive my critics of some of their favorite targets.

among those critics i must especially thank richard rorty, who has been egging me on for years to collect and publish these essays. ernest lepore generously gave up a week of his time to sort through and help order not only the essays in this volume but also those to appear in two subsequent volumes. i am indebted to ariela lazar

#### xviii introduction

who earlier gave me her discriminating advice on arranging my work, and to arpy khatchirian, who corrected the spelling, grammar, and thinking in many of the essays. peter momtchiloff of the oxford university press has been an encouraging, kindly, and forgiving editor; his assistance has made a task i found distasteful in prospect more endurable than i imagined possible.

a number of people, lectureships, universities, and other institutions have provided opportunities to try out many of the ideas in these essays. there has been welcome feedback from my students at berkeley, and from audiences who attended lecture series in mexico (1992), rome (1993), munich (the kant lectures, 1993), gerona (the ferrater mora lectures, 1994), leuven (the francqui lectures, 1994), and buenos aires (1995). finally, there were my jean nicod lectures (1995), delivered in caen and paris, from which i hope to develop a more unified and detailed account of the thoughts scattered through the present volume.

i am indebted to all those who added to, subtracted from, or modified my thoughts. i did not keep track of those who made especially trenchant suggestions, so any list i draw up will be shamefully incomplete. but it must certainly include rosario egidi, pascal engel, dagfinn f011esdal, olbeth hansberg, dieter henrich, pierre jacob, carol rovane, and those who came to the ten lectures and seminars in gerona, particularly w. v. quine, burton dreben, akeel bilgrami, ernest lepore, barry stroud, and bruce vermazen. marcia cavell not only attended many of the talks where i tried out versions of my ideas, but was an intellectual companion throughout these years, gently trying to temper my armchair speculations with a more empirically oriented, and psychoanalytically educated, outlook.

# 1 first person authority

when a speaker avers that he has a belief, hope, desire or intention, there is a presumption that he is not mistaken, a presumption that does not attach to his ascriptions of similar mental states to others. why should there be this asymmetry between attributions of attitudes to our present selves and attributions of the same attitudes to other selves? what accounts for the authority accorded first person present tense claims of this sort, and denied second or third person claims?

the point may be made, and the question asked, in the modality either of language or of epistemology. for if one can speak with special authority, the status of one's knowledge must somehow accord; while if one's knowledge shows some systematic difference, claims to know must reflect the difference. i assume therefore that if first person authority in speech can be explained, we will have done much, if not all, of what needs to be done to characterize and account for the epistemological facts.

the connection between the problem of first person authority and the traditional problem of other minds is obvious, but as i pose the former problem, there are two important differences. first person authority is the narrower problem, since i shall consider it only as it applies to propositional attitudes like belief, desire, intention; being pleased, astonished, afraid, or proud that something is the case; or knowing, remembering, noticing, or perceiving that something is the case. but i shall not discuss what are often taken to be central to the problem of other minds: pains and other sensations, and knowledge, memory, attention, and perception as directed to objects like people, streets, cities, comets, and other non-propositional entities. what holds for the propositional attitudes ought, it seems, to be relevant to sensations and the rest, but i do not explore the connections here.

all propositional attitudes exhibit first person authority, but in various degrees and kinds. belief and desire are relatively clear and simple examples, while intention, perception, memory, and knowledge are in one way or another more complex. thus in evaluating someone's claim to have noticed that the house is on fire, there are at least three things to consider: whether the house is on fire, whether the speaker believes the house is on fire, and how the fire caused the belief. with respect to the first, the speaker has no special authority; with respect to the second, he does; and with respect to the third, responsibility is mixed and complex. the question whether someone intends to lock the door by turning the key depends in part on whether he wants to lock the door and believes that turning the key will lock the door; and whether this belief and desire have caused, in the right way, a desire to turn the key, special authority attaches directly to claims about the desire and belief, less directly to claims about the necessary causal connection, these differences among the ways in which first person authority may apply to propositional attitudes are important and worth exploring. but in every case, first person authority is relevant, and it is the general case i wish to consider here. since in almost every instance, if not in all, first person authority rests at least partially on a belief component, i shall concentrate on the case of belief.

though there is first person authority with respect to beliefs and other propositional attitudes, error is possible; this follows from the fact that the attitudes are dispositions that manifest themselves in various ways, and over a span of time. error is possible; so is doubt. so we do not always have indubitable or certain knowledge of our own attitudes. nor are our claims about our own attitudes incorrigible. it is possible for the evidence available to others to overthrow self-judgements.

it comes closer to characterizing first person authority to note that the self-attributer does not normally base his claims on evidence or observation, nor does it normally make sense to ask the self-attributer why he believes he has the beliefs, desires, or intentions he claims to have. this feature of self-attributions was remarked by wittgenstein: 'what is the criterion for the redness of an image? for me, when it is someone else's image: what he says or does. for myself, when it is my image: nothing.'<sup>1</sup> most philosophers have

<sup>&</sup>lt;sup>1</sup> ludwig wittgenstein, philosophical investigations, §377.

followed wittgenstein in this, and have extended the criterion to the propositional attitudes, as we shall see.

this feature of first person authority, suggestive as it may be, does not help explain the authority. this is so partly because of the caveats—'normally' we do not make self-attributions on the basis of evidence, but sometimes we do; 'usually' it doesn't make sense to ask someone why he believes he has a certain belief or desire or intention, but sometimes it does. even in the exceptional cases, however, first person authority persists; even when a self-attribution is in doubt, or a challenge is proper, the person with the attitude speaks about it with special weight.

but the existence of exceptions is not the chief reason first person authority isn't explained by the fact that self-attributions are not based on evidence; the chief reason is simply that claims that are not based on evidence do not in general carry more authority than claims that are based on evidence, nor are they more apt to be correct.

contemporary philosophers who have discussed first person authority have made little attempt to answer the question why selfascriptions are privileged. it is long out of fashion to explain selfknowledge on the basis of introspection. and it is easy to see why, since this explanation leads only to the question why we should see any better when we inspect our own minds than when we inspect the minds of others.

a few philosophers have denied that the asymmetry exists; ryle is a sturdy example. in *the concept of mind* ryle suggests that what we take for 'privileged access' is due to nothing more than the fact that we are generally better placed to observe ourselves than others are. ryle writes, 'in principle, as distinct from practice, john doe's ways of finding out about john doe are the same as john doe's ways of finding out about richard roe'. he continues,

the differences are differences of degree, not of kind. the superiority of the speaker's knowledge of what he's doing over that of the listener does not indicate that he has privileged access to facts of a type inevitably inaccessible to the listener, but only that he is in a very good position to know what the listener is often in a very poor position to know. the turns taken by a man's conversation do not startle or perplex his wife as much as they had surprised and puzzled his fiancee, nor do close colleagues have to explain themselves to their new pupils.<sup>2</sup>

<sup>2</sup> gilbert ryle, the concept of mind, 156. 179.

i agree with ryle that any attempt to explain the asymmetry between first person present tense claims about attitudes, and other person or other tense claims, by reference to a special way of knowing or a special kind of knowledge must lead to a skeptical result. any such account must accept the asymmetry, but cannot explain it. but ryle neither accepts nor explains the asymmetry; he simply denies that it exists. since i think it is obvious that the asymmetry exists, i believe it is a mistake to argue from the absence of a special way of knowing or a special mode or kind of knowledge to the absence of special authority; instead, we should look for another source of the asymmetry.

aver at one time took a line similar to ryle's. in the concept of a *person* he emphasizes that first person ascriptions can be in error; and he allows that such ascriptions are privileged.<sup>3</sup> but when he comes to describe the authority of self-ascriptions, he compares it to the authority we sometimes allow an eyewitness when compared with secondhand reports, this analogy seems to me unsatisfactory for two reasons. first, it fails to tell us why a person is like an eyewitness with respect to his own mental states and events while others are not. and second, it does not suggest an accurate description of what first person authority is like. for first person attributions are not based on better evidence but often on no evidence at all, the authority of the eyewitness is at best based on inductive probabilities easily overridden in particular cases; an eyewitness is discredited and his evidence discounted if he is a notoriously unreliable observer, prejudiced, or myopic. but a person never loses his special claim to be right about his own attitudes, even when his claim is challenged or overturned.

joseph agassi has actually maintained that we know the mental states and events in other minds better than those in our own mind. he distinguishes privileged access from the commonsense truth that 'every person has access to some information available to that person alone, and it involves one's self, at least as an eye-witness'. he goes on:

the doctrine of privileged access is that i am the authority on all my own experiences . . . the thesis was refuted by freud (i know your dreams better than you), duhem (i know your methods of scientific discovery better than you), malinowski (i know your customs and habits better than you), and

<sup>3</sup> a. j. ayer, 'privacy'.

perception theorists (i can make you see things which are not there and describe your perceptions better than you can).

aside from freud's case, there is little here to threaten first person authority. freud's views, by extending the concepts of intention, belief, desire, and the rest to include the unconscious, do mean that with respect to some propositional attitudes a person loses direct authority. indeed, loss of authority is the main distinguishing feature of unconscious mental states. of course, the pre-freudian attitudes remain as subject as ever to first person authority. but more interesting is the fact that in psychoanalytic practice, recovery of authority over an attitude is often considered the only solid evidence that the attitude was there before being noninferentially appreciated by its holder. thus those cases of unconscious mental states that were unsystematically recognized to exist before freud are indirectly included in the scope of first person authority by psychoanalysis. so i do not think the existence of unconscious attitudes threatens the importance of first person authority.

i turn now to philosophers who have assumed that there is such a thing as first person authority, and have accepted wittgenstein's description of the difference between first and third person attributions.

strawson discusses first person authority in the context of trying to answer skepticism about other minds. according to strawson, if the skeptic understands his own question ('how does anyone know what is going on in someone else's mind?'), he knows the answer. for if the skeptic knows what a mind is, he knows it must be in a body, and that it has thoughts. he also knows that we attribute thoughts to others on the basis of observed behavior, but to ourselves without such a basis. strawson writes:

in order to *have* this type of concept [of a mental property], one must be both a self-ascriber and an other-ascriber of such predicates [which ascribe mental properties], and must see every other as a self-ascriber. in order to *understand* this type of concept, one must acknowledge that there is a kind of predicate which is unambiguously and adequately ascribable *both* on the basis of observation of the subject of the predicate *and* not on that basis, i.e. independently of observation of the subject.<sup>5</sup>

<sup>&</sup>lt;sup>4</sup> joseph agassi, *science in flux*, 120.

<sup>&</sup>lt;sup>5</sup> peter strawson, *individuals*, 108. (anita avramides has pointed out that strawson has contributed more to this issue than i allow. see her 'davidson and the new sceptical problem'.)

this cannot he deemed a satisfactory answer to the skeptic. for the skeptic will reply that though strawson may have correctly described the asymmetry between first and other person ascriptions of mental predicates, he has done nothing to explain it. in the absence of an explanation, the skeptic is surely justified in asking how we know that the description is correct. in particular, why should we think that a predicate that is sometimes applied on the basis of observation, and sometimes not, is unambiguous? this question, to which strawson has not addressed himself, is a major source of skepticism about knowledge of other minds. (on the apparent ambiguity of mental predicates, see essay 2.)

richard rorty has attempted an explanation. we are asked to imagine that originally self-ascriptions were made on the basis of the same sort of observation or behavioral evidence as other-ascriptions. it was then noticed that people could ascribe mental properties to themselves without making observations or using behavioral evidence, and that self-ascriptions turned out in the long run to provide better explanations of behavior than third person ascriptions. so it became a linguistic convention to treat self-ascriptions as privileged: 'it became a constraint on explanations of behavior that they should fit all reported thoughts or sensations into the overall account being offered'.<sup>6</sup>

this account is not meant to be taken seriously as a piece of folk anthropology, but it is meant to make it seem reasonable that we should treat self-ascriptions as having special authority. but the question remains: what reason has rorty given to show that self-ascriptions not based on evidence concern the same states and events as ascriptions of the same mental predicates based on observation or evidence? there is a difference in kind in the ways the two sorts of ascription are made, and how they explain behavior is different. what rorty describes as the discovery that self-ascriptions not based on evidence explain behavior better will be described by the skeptic as the fact that what is being ascribed is on every count apparently different.

it may come as a surprise to realize that the philosophers i am discussing have not really dealt with the ancient problem of skepticism concerning knowledge of other minds. but i think it is easy to

<sup>&</sup>lt;sup>6</sup> richard rorty, 'incorrigibility as the mark of the mental', 416. rorty's account is derived from wilfred sellars, 'empiricism and the philosophy of mind'. with respect to the point at issue, sellars's account does not differ from rorty's.

explain. historically the problem has been seen from either a cartesian or an empiricist point of view, and both venues assume that each person knows what is in his own mind. the problem has therefore seemed to be that of supplying a basis for knowledge of other minds (and, of course, the external world). philosophers now realize that part of understanding mental concepts (or predicates) consists in knowing what kind of observable behavior justifies the ascription of these concepts to others. but this answer to the skeptic does nothing to explain first person authority, or the asymmetry between self-ascriptions and other-ascriptions. we can still ask why we believe these two sorts of ascription pertain to the same subject matter. and this question is a good one, whether or not we recognize its traditional skeptical ancestry.

perhaps it should be pointed out that no concepts aside from those applying to sensations, propositional attitudes, and the positions of our limbs show the sort of asymmetry we are discussing. many concepts can be applied on the basis of multiple criteria, but no others are such that ascribers *must*, on particular occasions, use different criteria. if we are to explain this anomaly and avoid an invitation to skepticism, the explanation should point to a natural asymmetry between other observers and ourselves, an asymmetry not simply invented to solve the problem.

the first step towards a solution depends on becoming clear about the entities to which first person authority applies. william alston proposes this principle to characterize the special status of self-attributions: 'each person is so related to propositions ascribing current mental states to himself that it is logically impossible both for him to believe that such a proposition is true and not to be justified in holding this belief while no one else is so related to such propositions.'<sup>7</sup>

for this suggestion to be plausible, we must suppose that the proposition jones expresses by the sentence 'i believe wagner died happy' is the same proposition as the proposition smith expresses by the sentence 'jones believes wagner died happy.' this is, of course, a highly questionable supposition. once more, the epistemic contrast goes unexplained; and in the absence of an explanation, the question arises what *reason* one has in any particular case to believe that the proposition entertained by jones and smith *is* the same. given only

<sup>&</sup>lt;sup>7</sup> william alston,'varieties of privileged access', 235.

a description of an epistemic difference, the natural conclusion is that the propositions differ.

i turn, then, to a formulation of sidney shoemaker's, which makes explicit mention of language: 'among the incorrigible statements are statements about . . . mental events, e.g. . . . reports of thoughts . . . these are incorrigible in the sense that if a person sincerely asserts such a statement it does not make sense to suppose, and nothing could be accepted as showing, that he is mistaken; i.e. that what he says is false.'<sup>8</sup>

i shall ignore the incorrigibility condition and substitute something less strong—something that amounts to first person authority. (this is perhaps reasonable, since shoemaker is mainly concerned with sensations such as pain, while i am exclusively concerned with propositional attitudes.) what is important here is that shoemaker assigns the presumption of correctness not to a kind of knowledge, but to a class of utterances, this idea might lead to an explanation of first person authority if the class of utterances could be specified in syntactic terms. unfortunately it cannot. if shoemaker is right, a speaker who sincerely uses a certain sort of sentence must be presumed to be right in what he says. but of course this holds only if the speaker knows he is using the privileged sort of sentence; if he is not, he is misusing language. what would constitute a misuse here? above all, one wants to say, sincerely asserting a sentence one has no special authority to assert. perhaps so; but this is just to reiterate the uninformative and unexplained claim that it is a convention of language to treat self-ascriptions with special respect. seen from the point of view of the interpreter, this implies that he should interpret self-ascriptions in such a way as to make them true-or to assign a special priority to their truth. the point of view of the interpreter is the only one we can take, given shoemaker's principle, and this deprives the principle of independent application: our only reason for saying the speaker has special authority on occasion is that we are prepared to treat his utterance as a self-ascription. in other words, self-ascriptions have special authority: true; and that is where we began.

no satisfactory explanation of the asymmetry between first and other person attributions of attitudes has yet emerged. still, focusing on sentences and utterances rather than propositions or meanings is

sidney shoemaker, self-knowledge and self-identity, 215-16.

a step in a promising direction. the reason for this is relatively simple. as long as we pose the problem in terms of the kind of warrant or authority someone has with respect to claims about an agent's attitude to a proposition (or a sentence with a given interpretation), we seem constrained to account for differences by simply postulating different kinds or sources of information. alternatively, we may postulate different criteria of application for the key concepts or words ('believes that', 'intends to', 'wishes that', etc.). but these moves do no more than restate the problem, as we have seen, and thereby invite skepticism about knowledge of the minds of others (or of our own mind). but if we pose the problem in terms of relations between agents and utterances, we can avoid the impasse.

we now need to distinguish two related but different asymmetries. on the one hand, there is the familiar difference between self- and other-attributions of the same attitude to the same person: my claim that i believe wagner died happy and your claim that i believe wagner died happy, if these claims are put into words, we have the difficulty of deciding what pairs of utterances are suitably related in order to guarantee that the claims have the 'same content'. on the other hand, we may consider my utterance of the sentence 'i believe wagner died happy', and then contrast my warrant for thinking i have said something true, and your warrant for thinking i have said something true. these two asymmetries are of course connected since we are inclined to say your warrant for thinking i speak the truth when i say 'i believe wagner died happy' must be closely related to your warrant for thinking you would be speaking the truth if you said 'davidson believes wagner died happy'. but for reasons that will soon be evident, i shall deal with the second version of the asymmetry.

the question then comes to this: what explains the difference in the sort of assurance you have that i am right when i say 'i believe wagner died happy' and the sort of assurance i have? we know by now that it is no help to say i have access to a way of knowing about my own beliefs that you do not have; nor that we use different criteria in applying the concept of belief (or the word 'believes'). so let us simply consider a shorter utterance of mine: i utter the sentence 'wagner died happy'. clearly, if you or i or anyone knows that i hold this sentence true on this occasion of utterance, and she knows what i meant by this sentence on this occasion of utterance, then she knows what i believe—what belief i expressed.

it would once more make the account circular to explain the basic asymmetry by assuming an asymmetry in the assurance you and i have that i hold the sentence i have just uttered to be a true sentence. there must *be* such an asymmetry, of course, but it cannot be allowed to contribute to the desired explanation. but we can assume without prejudice that we both know, whatever the source or nature of our knowledge, that on this occasion i do hold the sentence i uttered to be true. similarly, it would beg the question to explain the basic asymmetry by appeal to some asymmetry in our knowledge of the fact that i know what my sentence, as uttered on this occasion, meant. so again, let us simply assume we both know this, whatever the source or character of our knowledge.

so far, then, we have not postulated or assumed any asymmetry at all. the assumptions are just these: you and i both know that i held the sentence 'wagner died happy' to be a true sentence when i uttered it; and that i knew what that sentence meant on the occasion of its utterance. and now there is this difference between us, which is what was to be explained: on these assumptions, i know what i believe, while you may not.

the difference follows, of course, from the fact that the assumption that i know what i mean necessarily gives me, but not you, knowledge of what belief i expressed by my utterance. it remains to show why there must be a presumption that speakers, but not their interpreters, are not wrong about what their words mean. the presumption is essential to the nature of interpretation—the process by which we understand the utterances of a speaker. this process cannot be the same for the utterer and for his hearers.

to put the matter in its simplest form: there can be no general guarantee that a hearer is correctly interpreting a speaker; however easily, automatically, unreflectively, and successfully a hearer understands a speaker, he is liable to serious error. in this special sense, he may always be regarded as interpreting a speaker. the speaker cannot, in the same way, interpret his own words. a hearer interprets (normally without thought or pause) on the basis of many clues: the actions and other words of the speaker, what he assumes about the education, birthplace, wit, and profession of the speaker, the relation of the speaker to objects near and far, and so forth. the speaker, though he must bear many of these things in mind when he speaks, since it is up to him to try to be understood, cannot wonder whether he generally means what he says.

#### *first person authority* 13

the contrast between the grounds a self-ascriber has for his selfascription, and the grounds an interpreter has for accepting that same ascription, would be stark if we were to assume that no question can arise concerning a speaker's interpretation of his own words. but of course it can, since what his words mean depends in part on the clues to interpretation he has given the interpreter, or other evidence that he justifiably believes the interpreter has. the speaker can be wrong about what his own words mean. this is one of the reasons first person authority is not infallible. but the possibility of error does not eliminate the asymmetry. the asymmetry rests on the fact that the interpreter must, while the speaker doesn't, rely on what, if it were made explicit, would be a difficult inference in interpreting the speaker.

neither speaker nor hearer knows in a special or mysterious way what the speaker's words mean; and both can be wrong. but there is a difference. the speaker, after bending whatever knowledge and craft he can to the task of saying what his words mean, cannot improve on the following sort of statement: 'my utterance of "wagner died happy" is true if and only if wagner died happy'. an interpreter has no reason to assume this will be *his* best way of stating the truth conditions of the speaker's utterance.

the best way to appreciate this difference is by imagining a situation in which two people who speak unrelated languages, and are ignorant of each other's languages, are left alone to learn to communicate. deciphering a new language is not like learning a first language, for a true beginner has neither the reasoning power nor the stock of concepts the participants in the imagined situation have to draw on. this does not, however, affect the point i wish to stress, since what my imagined interpreter can treat consciously as evidence is exactly what conditions the first learner to be a language user. let one of the imagined pair speak and the other try to understand, it will not matter whether the speaker speaks his 'native' tongue, since his past social situation is irrelevant. (i assume the speaker has no interest in training the hearer to cope with the speaker's original speech community.) the best the speaker can do is to be interpretable, that is, to use a finite supply of distinguishable sounds applied consistently to objects and situations he believes are apparent to his hearer. obviously the speaker may fail in this project from time to time; in that case we can say if we please that he does not know what his words mean. but it is equally obvious that the interpreter has nothing

to go on but the pattern of sounds the speaker exhibits in conjunction with further events (including, of course, further actions on the part of both speaker and interpreter). it makes no sense in this situation to wonder whether the speaker is generally getting things wrong. his behavior may simply not be interpretable. but if it is, then what his words mean is (generally) what he intends them to mean. since the 'language' he is speaking has no other hearers, the idea of the speaker misusing his language has no application. there is a presumption—an unavoidable presumption built into the nature of interpretation—that the speaker usually knows what he means. so there is a presumption that if he knows that he holds a sentence true, he knows what he believes.

## 2 knowing one's own mind

there is no secret about the nature of the evidence we use to decide what other people think: we observe their acts, read their letters, study their expressions, listen to their words, learn their histories, and note their relations to society, how we are able to assemble such material into a convincing picture of a mind is another matter; we know how to do it without necessarily knowing how we do it. sometimes i learn what i believe in much the same way someone else does, by noticing what i say and do. there may be times when this is my only access to my own thoughts. according to graham wallas, the little girl had the making of a poet in her who, being told to be sure of her meaning before she spoke, said 'how can i know what i think till i see what i say?'1 a similar thought was expressed by robert motherwell: 'i would say that most good painters don't know what they think until they paint it.'

gilbert ryle was with the poet and the painter all the way in this matter; he stoutly maintained that we know our own minds in exactly the same way we know the minds of others, by observing what we say, do, and paint. ryle was wrong. it is seldom the case that i need or appeal to evidence or observation in order to find out what i believe; normally i know what i think before i speak or act. even when i have evidence, i seldom make use of it. i can be wrong about my own thoughts, and so the appeal to what can be publicly determined is not irrelevant. but the possibility that one may be mistaken about one's own thoughts cannot defeat the overriding presumption that a person knows what he or she believes; in general, the belief that one has a thought is enough to justify that belief. but though this

1 graham wallas. the an of thought.

is true, and even obvious to most of us, the fact has, so far as i can see, no easy explanation. while it is clear enough, at least in outline, what we have to go on in trying to fathom the thoughts of others, it is obscure why, in our own case, we can so often know what we think without appeal to evidence or recourse to observation.

because we usually know what we believe (and desire and doubt and intend) without needing or using evidence (even when it is available), our sincere avowals concerning our present states of mind are not subject to the failings of conclusions based on evidence. thus sincere first person present-tense claims about thoughts, while neither infallible nor incorrigible, have an authority no second or third person claim, or first person other-tense claim, can have. to recognize this fact is not, however, to explain it.

since wittgenstein it has become routine to try to relieve worries about 'our knowledge of other minds' by remarking that it is an essential aspect of our use of certain mental predicates that we apply them to others on the basis of behavioral evidence but to ourselves without benefit of such aid. the remark is true, and when properly elaborated, it ought to answer someone who wonders how we can know the minds of others. but as a response to the skeptic, wittgenstein's insight (if it is wittgenstein's) should give little satisfaction, for first, it is a strange idea that claims made without evidential or observational support should be favored over claims with such support. of course, if evidence is not cited in support of a claim, the claim cannot be impugned by questioning the truth or relevance of the evidence. but these points hardly suffice to suggest that in general claims without evidential support are more trustworthy than those with the second, and chief, difficulty is this one would normally say that what counts as evidence for the application of a concept helps define the concept, or at least places constraints on its identification. if two concepts regularly depend for their application on different criteria or ranges of evidential support, they must be different concepts. so if what is apparently the same expression is sometimes correctly employed on the basis of a certain range of evidential support and sometimes on the basis of another range of evidential support (or none), the obvious conclusion would seem to be that the expression is ambiguous. why then should we suppose that a predicate like 'x believes that ras dashan is the highest mountain in ethiopia', which is applied sometimes on the basis of behavioral evidence and sometimes not, is unambiguous? if it is

ambiguous, then there is no reason to suppose it has the same meaning when applied to oneself that it has when applied to another. if we grant (as we should) that the necessarily public and interpersonal character of language guarantees that we often correctly apply these predicates to others, and that therefore we often do know what others think, then the question must be raised what grounds each of us has for thinking he knows what (in the same sense) he thinks. the wittgensteinian style of answer may solve the problem of other minds, but it creates a corresponding problem about knowledge of one's own mind. the correspondence is not quite complete, however, the original problem of other minds invited the question how one knows others have minds at all. the problem we now face must be put this way: i know what to look for in attributing thoughts to others. using quite different criteria (or none), i apply the same predicates to myself; so the skeptical question arises why i should think it is thoughts i am attributing to myself. but since the evidence i use in the case of others is open to the public, there is no reason why i shouldn't attribute thoughts to myself in the same way i do to others, in the mode of graham wallas, robert motherwell, and gilbert ryle. in other words, i don't, but i could, treat my own mental states in the same way i do those of others. no such strategy is available to someone who seeks the same sort of authority with respect to the thoughts of others as he apparently has in dealing with his own thoughts. so the asymmetry between the cases remains a problem, and it is first person authority that creates the problem.

i have suggested an answer to this problem in the first essay in this book. there i argued that attention to how we attribute thoughts and meanings to others would explain first person authority without inviting skeptical doubts. in recent years, however, some of the very facts about the attribution of attitudes on which i relied to defend first person authority have been employed to attack that authority: it has been argued, on what are thought to be new grounds, that while the methods of the third person interpreter determine what we usually deem to be the contents of an agent's mind, the contents so determined may be unknown to the agent. in the present essay i consider some of these arguments, and urge that they do not constitute a genuine threat to first person authority. the explanation i offered in my earlier essay of the asymmetry between first and other person attributions of attitudes seems to me if anything to be

strengthened by the new considerations, or those of them that seem valid.

it should be stressed again that the problem i am concerned with does not require that our beliefs about our own contemporary states of mind be infallible or incorrigible. we can and do make mistakes about what we believe, desire, approve, and intend; there is also the possibility of self-deceit. but such cases, though not infrequent, are not and could not be standard; i do not argue for this now, but take it as one of the facts to be explained.

setting aside, then, self-deception and other anomalous or borderline phenomena, the question is whether we can, without irrationality, inconsistency, or confusion, simply and straightforwardly think we have a belief we do not have, or think we do not have a belief we do have. a number of philosophers and philosophically minded psychologists have recently entertained views that entail or suggest that this could easily happen—indeed, that it must happen all the time.

the threat was there in russell's idea of propositions that could be known to be true even though they contained 'ingredients' with which the mind of the knower was not acquainted; and as the study of the *de re* attitudes evolved, the peril grew more acute.

but it was hilary putnam who pulled the plug. consider putnam's 1975 argument to show that meanings, as he put it, 'just ain't in the head'.<sup>2</sup> putnam argues persuasively that what words mean depends on more than 'what is in the head'. he tells a number of stories the moral of which is that aspects of the natural history of how someone learned the use of a word necessarily make a difference to what the word means. it seems to follow that two people might be in physically identical states, and yet mean different things by the same words.

the consequences are far-reaching. for if people can (usually) express their thoughts correctly in words, then their thoughts—their beliefs, desires, intentions, hopes, expectations—also must in part be identified by events and objects outside the person. if meanings ain't in the head, then neither, it would seem, are beliefs and desires and the rest.

since some may be a little weary of putnam's doppelganger on twin earth, let me tell my own science fiction story—if that is what

<sup>2</sup> hilary putnam, 'the meaning of "meaning" ', 227.

it is. my story avoids some irrelevant difficulties in putnam's story, though it introduces some new problems of its own.<sup>3</sup> (i'll come back to earth, and twin earth, a little later.) suppose lightning strikes a dead tree in a swamp; i am standing nearby. my body is reduced to its elements, while entirely by coincidence (and out of different molecules) the tree is turned into my physical replica. my replica, swampman, moves exactly as i did; according to its nature it departs the swamp, encounters and seems to recognize my friends, and appears to return their greetings in english. it moves into my house and seems to write articles on radical interpretation. no one can tell the difference.

but there *is* a difference. my replica can't recognize my friends; it can't recognize anything, since it never cognized anything in the first place. it can't know my friends' names (though of course it seems to); it can't remember my house. it can't mean what i do by the word 'house', for example, since the sound 'house' swampman makes was not learned in a context that would give it the right meaning—or any meaning at all. indeed, i don't see how my replica can be said to mean anything by the sounds it makes, nor to have any thoughts.

putnam might not go along with this last claim, for he says that if two people (or objects) are in relevantly similar physical states, it is 'absurd' to think their psychological states are 'one bit different'.<sup>4</sup> it would be a mistake to be sure that putnam and i disagree on this point, however, since it is not yet clear how the phrase 'psychological state' is being used.

putnam holds that many philosophers have wrongly assumed that psychological states like belief and knowing the meaning of a word both are (i) 'inner' in the sense that they do not presuppose the existence of any individual other than the subject to whom the state is ascribed, and are (ii) the very states which we normally identify and individuate as we do beliefs and the other propositional attitudes. since we normally identify and individuate mental states and meanings in

<sup>3</sup> i make no claim for originality here; stephen stich has used a very similar example in 'autonomous psychology and the belief-desire thesis'. 573 ff. i should emphasize that i am not suggesting that an object accidentally or artificially created could not think; the swampman simply needs time in which to acquire a causal history that would make sense of the claim that he is speaking of. remembering, identifying, or thinking of items in the world. (1 return to this point later.)

<sup>&</sup>lt;sup>4</sup> hilary putnam. the meaning of "meaning" ', 144.

terms partly of relations to objects and events other than the subject, putnam believes that (i) and (ii) come apart: in his opinion, no states can satisfy both conditions.

putnam calls psychological states satisfying condition (i) 'narrow'. he thinks of such states as solipsistic, and associates them with descartes's view of the mental. putnam may consider these states to be the only 'true' psychological states; in much of his paper he omits the qualifier 'narrow', despite the fact that narrow psychological states (so called) do not correspond to the propositional attitudes as normally identified. not everyone has been persuaded that there is an intelligible distinction to be drawn between narrow (or inner, or cartesian, or individualistic-all these terms are current) psychological states and psychological states identified (if any are) in terms of external facts (social or otherwise). thus john searle has claimed that our ordinary propositional attitudes satisfy condition (i), and so there is no need of states satisfying condition (ii), while tyler burge has denied that there are, in any interesting sense, propositional attitudes that satisfy condition (i).5 but there seems to be universal agreement that no states satisfy both conditions.

the thesis of this essay is that there is no reason to suppose that ordinary mental states do not satisfy both conditions (i) and (ii): i think such states are 'inner', in the sense of being identical with states of the body, and so identifiable without reference to objects or events outside the body; they are at the same time 'non-individualistic' in the sense that they can be, and usually are, identified in part by their causal relations to events and objects outside the subject whose states they are. a corollary of this thesis will turn out to be that, contrary to what is often assumed, first person authority can without contradiction apply to states that are regularly identified by their relations to events and objects outside the person.

i begin with the corollary. why is it natural to assume that states that satisfy condition (ii) may not be known to the person who is in those states?

now i must talk about putnam's twin earth. he asks us to imagine two people exactly alike physically and (therefore) alike with respect to all 'narrow' psychological states. one of the two people, an inhabitant of earth, has learned to use the word 'water' by being shown water, reading and hearing about it, etc. the other, an inhabitant of

john searle, intentionality; and tyler burge. 'individualism and psychology'.

twin earth, has learned to use the word 'water' under conditions not observably different, but the substance to which she has been exposed is not water but a lookalike substance we may call 'twater'. under the circumstances, putnam claims, the first speaker refers to water when she uses the word 'water'; her twin refers to twater when *she* uses the word 'water'. so we seem to have a case where 'narrow' psychological states are identical, and yet the speakers mean different things by the same word.

how about the thoughts of these two speakers? the first says to herself, when facing a glass of water, 'here's a glass of water'; the second mutters exactly the same sounds to herself when facing a glass of twater. each speaks the truth, since their words mean different things. and since each is sincere, it is natural to suppose they believe different things, the first believing there is a glass of water in front of her, the second believing there is a glass of twater in front of her. but do they know what they believe? if the meanings of their words, and thus the beliefs expressed by using those words, are partly determined by external factors about which the agents are ignorant, their beliefs and meanings are not narrow in putnam's sense. there is therefore nothing on the basis of which either speaker can tell which state she is in, for there is no internal or external clue to the difference available, we ought, it seems, to conclude that neither speaker knows what she means or thinks. the conclusion has been drawn explicitly by a number of philosophers, among them putnam. putnam declares that he 'totally abandons the idea that if there is a difference in meaning . . . then there must be some difference in our concepts (or in our psychological state)'. what determines meaning and extension 'is not, in general, fully known to the speaker'.<sup>6</sup> here 'psychological state' means *narrow* psychological state, and it is assumed that only such states are 'fully known'. jerry fodor believes that ordinary propositional attitudes are (pretty nearly) 'in the head', but he agrees with putnam that *if* propositional attitudes were partly identified by factors outside the agent, they would not be in the head, and would not necessarily be known to the agent.<sup>7</sup> john searle also, though his reasons are not fodor's, holds that meanings are in the head ('there is nowhere else for them to be'),

<sup>&</sup>lt;sup>6</sup> hilary putnam. 'the meaning of "meaning" ' pp. 164-5.

<sup>&</sup>lt;sup>7</sup> jerry fodor, 'cognitive science and the twin earth problem', 103. also see his 'methodological solipsism considered as a research strategy in cognitive psychology', p. viii.

but seems to accept the inference that if this were not the case, first person authority would be lost.8 perhaps the plainest statement of the position appears in andrew woodfield's introduction to a book of essays on the objects of thought. referring to the claim that the contents of the mind are often determined by facts external to and perhaps unknown to the person whose mind it is, he says: 'because the external relation is not determined subjectively, the subject is not authoritative about that. a third person might well be in a better position than the subject to know which object the subject is thinking about, hence be better placed to know which thought it was.<sup>9</sup> those who accept the thesis that the contents of propositional attitudes are partly identified in terms of external factors seem to have a problem similar to the problem of the skeptic who finds that we may be altogether mistaken about the 'outside' world. in the present case, ordinary skepticism of the senses is avoided by supposing that the world itself more or less correctly determines the contents of thoughts about the world. (the speaker who thinks it is water is probably right, for he learned the use of the word 'water' in a watery environment; the speaker who thinks it is twater is probably right, for he learned the word 'water' in a twatery environment.) but skepticism is not defeated; it is only displaced onto knowledge of our own minds, our ordinary beliefs about the external world are (on this view) directed onto the world, but we don't know what we believe.

there is, of course, a difference between water and twater, and it can be discovered by normal means, whether it is discovered or not. so a person might find out what he believes by discovering the difference between water and twater, and finding out enough about his own relations to both to determine which one his talk and beliefs are about. the skeptical conclusion we seem to have reached concerns the extent of first person authority: it is far more limited than we supposed. our beliefs about the world are mostly true, but we may easily be wrong about what we think. it is a transposed image of cartesian skepticism.

those who hold that the contents of our thoughts and the meanings of our words are often fixed by factors of which we are ignorant have not been much concerned with the apparent consequence of their views which i have been emphasizing. they have, of course,

<sup>&</sup>lt;sup>8</sup> john searle, *intentionality*, ch. 8.

<sup>&</sup>lt;sup>9</sup> andrew woodfield 'introduction', in woodfield (ed.), thought and object, p. viii.
realized that if they were right, the cartesian idea that the one thing we can be certain of is the contents of our own minds, and the fregean notion of meanings fully 'grasped', must be wrong. but they have not made much of an attempt, so far as i know, to resolve the seeming conflict between their views and the strong intuition that first person authority exists.

one reason for the lack of concern may be that some seem to see the problem as confined to a fairly limited range of cases, cases where concepts or words latch on to objects that are picked out or referred to using proper names, indexicals, and words for natural kinds, others, though, argue that the ties between language and thought, on the one hand, and external affairs, on the other, are so pervasive that no aspect of thought as usually conceived is untouched. in this vein daniel dennett remarks that 'one must be richly informed about, intimately connected with, the world at large, its occupants and properties, in order to be said with any propriety to have beliefs'.<sup>10</sup> he goes on to claim that the identification of all beliefs is infected by the outside, non-subjective factors that are recognized to operate in the sort of case we have been discussing. burge also emphasizes the extent to which our beliefs are affected by external factors, though for reasons he does not explain, he apparently does not view this as a threat to first person authority."

the subject has taken a disquieting turn. at one time behaviorism was invoked to show how it was possible for one person to know what was in another's mind; behaviorism was then rejected, in part because it could not explain one of the most obvious aspects of mental states: the fact that they are in general known to the person who has them without appeal to behavioristic evidence. the recent fashion, though not strictly behavioristic, once more identifies mental states partly in terms of social and other external factors, thus making them to that extent publicly discoverable. but at the same time it reinstates the problem of accounting for first person authority.

those who are convinced of the external dimension of the contents of thoughts as ordinarily identified and individuated have reacted in different ways. one response has been to make a distinction between the contents of the mind as subjectively and internally determined, on

<sup>&</sup>lt;sup>10</sup> daniel dennett, 'beyond belief, 76.

<sup>&</sup>lt;sup>11</sup> tyler burge, 'other bodies', 'individualism and the mental', 'two thought experiments reviewed'. 284-93.

the one hand, and ordinary beliefs, desires, and intentions, as we normally attribute them on the basis of social and other outward connections, on the other. this is clearly the trend of putnam's argument (although the word 'water' has different meanings, and is used to express different beliefs when it is used to refer to water and to twater, people using the word for these different purposes may be in 'the same psychological state'). jerry fodor accepts the distinction for certain purposes, but argues that psychology should adopt the stance of 'methodological solipsism'—that is, it should deal exclusively with inner states, the truly subjective psychological states that owe nothing to their relations to the outside world.<sup>12</sup>

stephen stich makes essentially the same distinction, but draws a sterner moral: where fodor thinks we merely need to tinker a bit with propositional attitudes as usually conceived to separate out the purely subjective element, stich holds that psychological states as we now think of them belong to a crude and confused 'folk psychology' which must be replaced by a yet to be invented 'cognitive science'. the subtitle of his recent book is the case against belief'.<sup>13</sup>

clearly those who draw such a distinction have insured that the problem of first person authority, at least as i have posed it, cannot be solved. for the problem i have set is how to explain the asymmetry between the way in which a person knows about his contemporary mental states and the way in which others know about them. the mental states in question are beliefs, desires, intentions, and so on, as ordinarily conceived. those who accept something like putnam's distinction do not even try to explain first person authority with respect to these states; if there is first person authority at all, it attaches to quite different states. (in stich's case, it is not obvious that it can attach to anything.)

i think putnam, burge, dennett, fodor, stich, and others are right in calling attention to the fact that ordinary mental states, at least the propositional attitudes, are partly identified by relations to society and the rest of the environment, relations which may in some respects not be known to the person in those states. they are also right, in my opinion, in holding that for this reason (if for no other), the concepts of 'folk psychology' cannot be incorporated into a

<sup>&</sup>lt;sup>12</sup> jerry fodor. 'methodological solipsism considered as a research strategy in cognitive psychology'.

<sup>&</sup>lt;sup>3</sup> stephen stich, from folk psychology to cognitive science.

coherent and comprehensive system of laws of the sort for which physics strives. these concepts are part of a commonsense theory for describing, interpreting, and explaining human behavior which is a bit freestyle, but (so i think) indispensable. i can imagine a science concerned with how people think and act purged of 'folk psychology', but i cannot think in what its interest would consist. this is not, however, the topic of this essay.

i am here concerned with the puzzling discovery that we apparently do not know what we think—at least in the way we think we do. this is a real puzzle if, like me, you believe it is true that external factors partly determine the contents of thoughts, and also believe that in general we do know, and in a manner others do not, what we think. the problem arises because admitting the identifying and individuating role of external factors seems to lead to the conclusion that our thoughts may not be known to us.

but does this conclusion follow? the answer depends on the way in which one thinks the identification of mental contents depends on external factors.

the conclusion does follow, for example, for any theory that holds that propositional attitudes are identified by objects (such as propositions, tokens of propositions, or representations) that are in or 'before' the mind, and that contain or incorporate (as 'ingredients') objects or events outside the agent; for it is obvious that everyone is ignorant of endless features of every external object. that the conclusion follows from these assumptions is generally conceded.<sup>14</sup> however, for reasons i shall mention below, i reject the assumptions on which the conclusion is in this case based.

tyler burge has suggested that there is another way in which external factors enter into the determination of the contents of speech and thought. one of his 'thought experiments' happens pretty well to fit me. until recently i believed arthritis was an inflammation of the joints caused by calcium deposits; i did not know that any inflammation of the joints, for example gout, also counted as arthritis. so when a doctor told me (falsely as it turned out) that i had gout, i believed i had gout but i did not believe i had arthritis. at this point burge asks us to imagine a world in which i was physically the same but in which the word 'arthritis' happened actually to apply only to inflammation of the joints caused by calcium deposits. then the

<sup>&</sup>lt;sup>14</sup> see e.g. gareth evans. the varieties of reference, 45. 199. 201.

sentence 'gout is not a form of arthritis' would have been true, not false, and the belief that i expressed by this sentence would not have been the false belief that gout is not a form of arthritis but a true belief about some disease other than arthritis. yet in the imagined world all my physical states, my 'internal qualitative experiences', my behavior and dispositions to behave, would have been the same as they are in this world. my *belief would* have changed, but i would have no reason to suppose that it had, and so could not be said to know what i believed.

burge stresses the fact that his argument depends on 'the possibility of someone's having a propositional attitude despite an incomplete mastery of some notion in its content . . . *if* the thought experiment is to work, one must at some stage find the subject believing for having some attitude characterized by) a content, despite an incomplete understanding or misapplication'.<sup>15</sup>

it seems to follow that if burge is right, then whenever a person is wrong, confused, or partially misinformed, about the public meaning of a word, he is wrong, confused, or partially misinformed about any of his beliefs that are (or would be?) expressed by using that word. since such 'partial understanding' is 'common or even normal in the case of a large number of expressions in our vocabularies' according to burge, it must be equally common or normal for us to be wrong about what we believe (and, of course, fear, hope for, wish were the case, doubt, and so on).

burge apparently accepts this conclusion; at least so i interpret his denial that 'full understanding of a content is in general a necessary condition for believing the content'. he explicitly rejects 'the old model according to which a person must be directly acquainted with, or must immediately apprehend, the contents of his thoughts ... a person's thought *content* is not fixed by what goes on in him, or by what is accessible to him simply by careful reflection'.<sup>16</sup>

i am uncertain how to understand these claims, since i am uncertain how seriously to take the talk of 'direct acquaintance' with, and of 'immediately apprehending', a content. but in any case i am convinced that if what we mean and think is determined by the linguistic habits of those around us in the way burge believes they are, then first person authority is very seriously compromised. since the degree and character of the compromise seem to me incompati-

tyler burge, 'individualism and the mental'. 83.

ibid. 90, 102. 104.

hie with what we know about the kind of knowledge we have of our own minds, i must reject some premise of surge's. i agree that what i mean and think is not 'fixed' (exclusively) by what goes on in me, so what i must reject is surge's account of how social and other external factors control the contents of a person's mind.

for a number of reasons, i am inclined to discount the importance of the features of our attributions of attitudes to which burge points. suppose that i, who think the word 'arthritis' applies to inflammation of the joints only if caused by calcium deposits, and my friend arthur, who knows better, both sincerely utter to smith the words 'carl has arthritis' according to burge, if other things are more or less equal (arthur and i are both generally competent speakers of english, both have often applied the word 'arthritis' to genuine cases of arthritis, etc.), then our words on this occasion mean the same thing, arthur and i mean the same thing by our words, and we express the same belief. my error about the dictionary meaning of the word (or about what arthritis is) makes no difference to what i meant or thought on this occasion. surge's evidence for this claim seems to rest on his conviction that this is what anyone (unspoiled by philosophy) would report about arthur and me. i doubt that burge is right about this, but even if he is, i don't think it proves his claim. ordinary attributions of meanings and attitudes rest on vast and vague assumptions about what is and is not shared (linguistically and otherwise) by the attributer, the person to whom the attribution is made, and the attributer's intended audience. when some of these assumptions prove false, we may alter the words we use to make the report, often in substantial ways, when nothing much hinges on it, we tend to choose the lazy way: we take someone at his word, even if this does not guite reflect some aspect of the speaker's thought or meaning. but this is not because we are bound (outside of a law court, anyway) to be legalistic about it. and often we aren't. if smith (unspoiled by philosophy) reports to still another party (perhaps a distant doctor attempting a diagnosis on the basis of a telephone report) that arthur and i both have said, and believe, that carl has arthritis, he may actively mislead *his* hearer, if this danger were to arise, smith, alert to the facts, would not simply say 'arthur and davidson both believe carl has arthritis'; he would add something like 'but davidson thinks arthritis must be caused by calcium deposits'. the need to make this addition i take to show that the simple attribution was not quite right; there was a relevant difference in the thoughts arthur and i expressed when we said 'carl

has arthritis'. burge does not have to be budged by this argument, of course, since he can insist that the report is literally correct, but could, like any report, be misleading. i think, on the other hand, that this reply would overlook the extent to which the contents of one belief necessarily depend on the contents of others. thoughts are not independent atoms, and so there can be no simple, rigid, rule for the correct attribution of a single thought.<sup>17</sup>

though i reject burge's insistence that we are bound to give a person's words the meaning they have in his linguistic community, and to interpret his propositional attitudes on the same basis, i think there is a somewhat different, but very important, sense in which social factors do control what a speaker can mean by his words. if a speaker wishes to be understood, he must intend his words to be interpreted in a certain way, and so must intend to provide his audience with the clues they need to arrive at the intended interpretation. this holds whether the hearer is sophisticated in the use of a language the speaker knows or is the learner of a first language. it is the requirement of learnability, interpretability, that provides the irreducible social factor, and that shows why someone can't mean something by his words that can't be correctly deciphered by another. (burge seems to make this point himself in a later paper.<sup>18</sup>)

now i would like to return to putnam's twin earth example, which does not depend on the idea that social linguistic usage dictates (under

this has nothing (directly) to do with what they meant by their words, nor what they believed, (b) as good citizens and parents we want to encourage practices that enhance the chances for communication; using words as we think others do may enhance communica tion. this thought (whether or not justified) may help explain why some people tend to attribute meanings and beliefs in a legalistic way; they hope to encourage conformity, (c) a speaker who wishes to be understood must intend his words to be interpreted (and hence interpretable) along certain lines; this intention may be served by using words as others do

(though often this is not the case). similarly, a hearer who wishes to understand a speaker must intend to interpret the speaker's words as the speaker intended (whether or not the interpretation is 'standard'). these reciprocal intentions become morally important in endless situations which have no necessary connection with the determination of what someone had in mind.

<sup>8</sup> tyler burge, two thought experiments reviewed', 289.

<sup>&</sup>lt;sup>17</sup> burge suggests that the reason we normally take a person to mean hy his words what others in his linguistic community mean, whether or not the speaker knows what others mean, is that 'people are frequently held, and hold themselves, to the standards of the community when misuse or misunderstanding are at issue.' he also says such cases 'depend on a certain responsibility to communal practice' (ibid. 90). i don't doubt the phenomenon, but doubt its bearing on what it is supposed to show, (*a*) it is often reason able to hold people responsible for knowing what their words mean; in such cases we may treat them as committed to positions they did not know or believe they were committed to.

#### knowing one's own mind 29

more or less standard conditions) what speakers mean by their words, nor, of course, what their (narrow) psychological states are. i am, as i said, persuaded that putnam is right; what our words mean is fixed in part by the circumstances in which we learned, and used, the words. putnam's single example (water) is not enough, perhaps, to nail down this point, since it is possible to insist that 'water' doesn't apply just to stuff with the same molecular structure as water but also to stuff enough like water in structure to be odorless, potable, to support swimming and sailing, etc. (i realize that this remark, like many others in this piece, may show that i don't know a rigid designator when i see one. i don't.) the issue does not depend on such special cases, nor on how we do or should resolve them, the issue depends simply on how the basic connection between words and things, or thoughts and things, is established. i hold, along with burge and putnam if i understand them, that it is established by causal interactions between people and parts and aspects of the world. the dispositions to react differentially to objects and events thus set up are central to the correct interpretation of a person's thoughts and speech. if this were not the case, we would have no way of discovering what others think, or what they mean by their words. the principle is as simple and obvious as this: a sentence someone is inspired (caused) to hold true by and only by sightings of the moon is apt to mean something like 'there's the moon'; the thought expressed is apt to be that the moon is there; the thought inspired by and only by sightings of the moon is apt to be the thought that the moon is there, apt to be, allowing for intelligible error, secondhand reports, and so on. not that all words and sentences are this directly conditioned to what they are about; we can perfectly well learn to use the word 'moon' without ever seeing it. the claim is that all thought and language must have a foundation in such direct historical connections, and these connections constrain the interpretation of thoughts and speech. perhaps i should stress that the arguments for this claim do not rest on intuitions concerning what we would say if certain counterfactuals were true. no science fiction or thought experiments are required.<sup>19</sup>

<sup>&</sup>lt;sup>19</sup> burge has described thought experiments which do not involve language at all: one of these experiments prompts him to claim that someone brought up in an environment without aluminum could not have 'aluminum thoughts' ('individualism and psychology', 5.) burge does not say why he thinks this, but it is by no means obvious that counterfactual assumptions are needed to make the point. in any case, the new thought experiments seem to rest on intuitions quite different from the intuitions invoked in 'individualism and the mental': it is not clear how social norms feature in the new experiments, and the linguistic habits of the community are apparently irrelevant. at this point it may be that burge's position is close to mine.

i agree with putnam and burge, then, that (as burge puts it): 'the intentional content of ordinary propositional attitudes . . . cannot he accounted for in terms of physical, phenomenal, causal-functional, computational, or syntactical states or processes that are specified nonintentionally and are defined purely on the individual in isolation from his physical and social environment'.<sup>20</sup> the question remains whether this fact is a threat to first person authority, as burge seems to think, and putnam and others certainly think. i have rejected one of burge's arguments which, if it were right, would pose such a threat. but there is the position described in the previous paragraph, which i hold whether or not others do, since i think this much 'externalism' is required to explain how language can be learned, and how words and attitudes can be identified by an interpreter.

why does putnam think that if the reference of a word is (sometimes) fixed by the natural history of how the word was acquired, a user of the word may lose first person authority? putnam claims (correctly, in my view) that two people can be in all relevant physical (chemical, physiological, etc.) respects the same and yet mean different things by their words and have different propositional attitudes (as these are normally identified). the differences are due to environmental differences about which the two agents may, in some respects, be ignorant, why, under these circumstances, should we suppose these agents may not know what they mean and think? talking with them will not easily show this. as we have noted, each, when faced with a glass of water or twater, says honestly, 'here's a glass of water'. if they are in their home environments, each is right; if they have switched earths, each is wrong. if we ask each one what he means by the word 'water', he gives the right answer, using the same words, of course. if we ask each one what he believes, he gives the right answer, these answers are right because, though verbally identical, they must be interpreted differently. and what is it that they do not know (in the usual authoritative way) about their own states? as we have seen, putnam distinguishes the states we have just been discussing from 'narrow' psychological states which do not presuppose the existence of any individual other than the subject in that state. we may now start to wonder why putnam is interested in narrow psychological states. part of the answer is, of course, that it is these states that he thinks have the 'cartesian' property of being

<sup>20</sup> two thought experiments reviewed'. 288.

known in a special way by the person who is in them. (the other part of the answer has to do with constructing a 'scientific psychology'; this does not concern us here.)

the reasoning depends, i think, on two largely unquestioned assumptions. these are:

- 26. if a thought is identified by a relation to something outside the head, it isn't wholly in the head. (it ain't in the head.)
- 27. if a thought isn't wholly in the head, it can't be 'grasped' by the mind in the way required by first person authority.

that this is putnam's reasoning is suggested by his claim that if two heads are the same, narrow psychological states must be the same. thus if we suppose two people are 'molecule for molecule' the same ('in the sense in which two neckties can be "identical" '; you may add, if you wish, that each of the two people 'thinks the same verbalized thoughts . . . has the same sense data, the same dispositions, etc.'), then 'it is absurd to think [one] psychological state is one bit different from' the other. these are, of course, narrow psychological states, not the ones we normally attribute, which ain't in the head.<sup>21</sup>

it is not easy to say in exactly what way the verbalized thoughts, sense data, and dispositions can be identical without reverting to the neckties, so let us revert. then the idea is this: the narrow psychological states of two people are identical when their physical states cannot be distinguished. there would be no point in disputing this, since narrow psychological states are putnam's to define; what i wish to question is assumption (1) above, which led to the conclusion that ordinary propositional attitudes aren't in the head, and that therefore first person authority doesn't apply to them.

it should be clear that it doesn't follow, simply from the fact that meanings are identified in part by relations to objects outside the head, that meanings aren't in the head. to suppose this would be as bad as to argue that because my being sunburned presupposes the existence of the sun, my sunburn isn't a condition of my skin. my sunburned skin may be indistinguishable from someone else's skin that achieved its burn by other means (our skins may be identical in the 'necktie sense'); yet one of us is really sunburned and the other not. this is enough to show that an appreciation of the external

<sup>21</sup> 'the meaning of "meaning". 227.

factors that enter into our common ways of identifying mental states does not discredit an identity theory of the mental and the physical. andrew wood field seems to think it does. he writes: 'no *de re* state about an object that is external to the person's brain can possibly be identical with a state of that brain, since no brain state presupposes the existence of an external object.'<sup>22</sup> individual states and events don't *conceptually* presuppose anything in themselves; some of their *descriptions* may, however. my paternal grandfather didn't presuppose me, but if someone can be described as my paternal grandfather, several people besides my grandfather, including me, must exist. burge may make a similar mistake in the following passage:

no occurrence of a thought... could have a different content and be the very same token event ... then ... a person's thought event is not *identical* with any event in him that is described by physiology, biology, chemistry, or physics. for let *b* be any given event described in terms of one of the physical sciences that occurs in the subject while he thinks the relevant thought. let 'b' be such that it denotes the same physical event occurring in the subject in our counterfactual situation ... *b* need not be affected by counterfactual differences [that do not change the contents of the thought event]. thus ... *b* [the physical event] is not identical with the subject's occurrent thought.<sup>23</sup>

burge does not claim to have established the premise of this argument, and so not its conclusion. but he holds that the denial of the premise is 'intuitively very implausible'. he goes on, 'materialist identity theories have schooled the imagination to picture the content of a mental event as varying while the event remains fixed. but whether such imaginings are possible fact or just philosophical fancy is a separate question'. it is because he thinks the denial of the premise to be very improbable that he holds that 'materialist identity theories' are themselves 'rendered implausible by the non-individualistic thought experiments'.<sup>24</sup>

i accept burge's premise; i think its denial not merely implausible but absurd. if two mental events have different contents, they are surely different events. what i take burge's and putnam's imagined cases to show (and what i think the swampman example shows more directly) is that people who are in all relevant physical respects similar (or 'identical' in the necktie sense) can differ in what they mean

<sup>&</sup>lt;sup>2</sup> andrew woodfield, in thought and object, p. viii.

<sup>&</sup>lt;sup>23</sup> 'individualism and the mental', 111.

<sup>&</sup>lt;sup>24</sup> 'individualism and psychology', 15 n. 7. cf. 'individualism and the mental', 11

or think, just as they can differ in being grandfathers or being sunburned. but of course there is *something* different about them, even in the physical world; their causal histories are different, and they are discrete physical objects.

i conclude that the mere fact that ordinary mental states and events are individuated in terms of relations to the outside world has no tendency to discredit mental-physical identity theories as such. in conjunction with a number of further (plausible) assumptions, the 'externalism' of certain mental states and events can be used, i think, to discredit type-type identity theories; but if anything it supports token-token identity theories. (i see no good reason for calling all identity theories 'materialist'; if some mental events are physical events, this makes them no more physical than mental. identity is a symmetrical relation.)

putnam and woodfield are wrong, then, in claiming that it is 'absurd' to think two people could be physically identical (in the 'necktie sense') and yet differ in their ordinary psychological states. burge, unless he is willing to make far stronger play than he has with essentialist assumptions, is wrong in thinking he has shown all identity theories implausible. we are therefore free to hold that people can be in all relevant physical respects identical (identical in the 'necktie sense') while differing psychologically: this is in fact the position of 'anomalous monism', for which i have argued elsewhere.<sup>25</sup>

one obstacle to nonevidential knowledge of our own ordinary propositional attitudes has now been removed. for if ordinary beliefs and the other attitudes can be 'in the head' even though they are identified as the attitudes they are partly in terms of what is not in the head, then the threat to first person authority cannot come simply from the fact that external factors are relevant to the identification of the attitudes.

but an apparent difficulty remains. true, my sunburn, though describable as such only in relation to the sun, is identical with a condition of my skin which can (i assume) be described without reference to such external factors. still, if, as a scientist skilled in all the physical sciences, i have access only to my skin, and am denied knowledge of the history of its condition, then by hypothesis there is no way for me to tell that i am sunburned. perhaps, then, someone

<sup>&</sup>lt;sup>25</sup> 'mental events', in donald davidson. essays on actions and events.

has first person authority with respect to the contents of his mind only as those contents can be described or discovered without reference to external factors. in so far as the contents are identified in terms of external factors, first person authority necessarily lapses. 1 can tell by examining my skin what my private or 'narrow' condition is, but nothing i can learn in this restricted realm will tell me that 1 am sunburned. the difference between referring to and thinking of water and referring to and thinking of twater is like the difference between being sunburned and one's skin being in exactly the same condition through another cause. the semantic difference lies in the outside world, beyond the reach of subjective or sublunar knowledge. so the argument might run.

this analogy, between the limited view of the skin doctor and the tunnel vision of the mind's eye, is fundamentally flawed. it depends for its appeal on a faulty picture of the mind, a picture which those who have been attacking the subjective character of ordinary psychological states share with those they attack. if we can bring ourselves to give up this picture, first person authority will no longer been seen as a problem; indeed, it will turn out that first person authority is dependent on, and explained by, the social and public factors that were supposed to undermine that authority.

there is a picture of the mind which has become so ingrained in our philosophical tradition that it is almost impossible to escape its influence even when its faults are recognized and repudiated. in one crude, but familiar, version it goes like this: the mind is a theater in which the conscious self watches a passing show (the shadows on the wall). the show consists of 'appearances', sense data, qualia, what is 'given' in experience. what appear on the stage are not the ordinary objects in the world that the outer eye registers and the heart loves, but their purported representatives. whatever we know about the world outside depends on what we can glean from the inner clues.

the difficulty that has been apparent from the start with this description of the mental is to see how it is possible to beat a reliable track from the inside to the outside. another conspicuous, though perhaps less appreciated, difficulty is to locate the self in the picture. for the self seems on the one hand to include theater, stage, actors, and audience; on the other hand, what is known and registered pertains to the audience alone. this second problem could be as well stated as the problem of the location of the objects of the mind: are they *in* the mind, or simply viewed *by* it?

i am not now concerned with such (now largely disavowed) objects of the mind as sense data, but with their judgemental cousins, the supposed objects of the propositional attitudes, whether thought of as propositions, tokens of propositions, representations, or fragments of 'mentalese'. the central idea i wish to attack is that these are entities that the mind can 'entertain', 'grasp', 'have before it', or be 'acquainted' with. (these metaphors are probably instructive: voyeurs merely want to have representations before the mind's eye, while the more aggressive grasp them; the english may be merely acquainted with the contents of the mind, while more friendly types will actually entertain them.)

it is easy to see how the discovery that external facts enter into the indi viduation of states of mind disturbs the picture of the mind i have been describing. for if to be in a state of mind is for the mind to be in some relation to an object, then whatever helps determine what object it is must equally be grasped if the mind is to know what state it is in. this is particularly evident if an external object is an 'ingredient' in the object before the mind. but in either case, the person who is in the state of mind may not know what state of mind he is in.

it is at this point that the concept of the subjective—of a state of mind—seems to come apart. on the one hand, there are the true inner states, with respect to which the mind retains its authority; on the other hand, there are the ordinary states of belief, desire, intention, and meaning, which are polluted by their necessary connections with the social and public world.

in analogy, there is the problem of the sunburn expert who cannot tell by inspecting the skin whether it is a case of sunburn or merely an 'identical' condition with another cause. we can solve the sunburn problem by distinguishing between sunburn and sunnishburn; sunnishburn is just like sunburn except that the sun may or may not be involved. the expert can spot a case of sunnishburn just by looking, but not a case of sunburn. this solution works because skin conditions, unlike objects of the mind, are not required to be such that there be a special someone who can tell, just by looking, whether or not the condition obtains.

the solution in the case of mental states is different, and simpler; it is to get rid of the metaphor of objects before the mind. most of us long ago gave up the idea of perceptions, sense data, the flow of experience, as things 'given' to the mind; we should treat prepositional objects in the same way. of course people have beliefs,

wishes, doubts, and so forth; but to allow this is not to suggest that beliefs, wishes, and doubts are *entities* in or before the mind, or that being in such states requires there to be corresponding mental objects.

this has been said before, in various tones of voice, but for different reasons. ontological scruples, for example, are no part of my interest. we will always need an infinite supply of objects to help describe and identify attitudes like belief; i am not suggesting for a moment that belief sentences, and sentences that attribute the other attitudes, are not relational in nature. what i am suggesting is that the objects to which we relate people in order to describe their attitudes need not in any sense be *psychological* objects, objects to be grasped, known, or entertained by the person whose attitudes are described.

this point, too, is familiar; quine makes it when he suggests that we may use our own sentences to keep track of the thoughts of people who do not know our language. quine's interest is semantical, and he says nothing in this context about the epistemological and psychological aspects of the attitudes. we need to bring these various concerns together. sentences about the attitudes are relational; for *semantic* reasons there must therefore be objects to which to relate those who have attitudes. but having an attitude is not having an entity before the mind; for compelling *psychological* and *epistemological* reasons we should deny that there are objects of the mind.

the source of the trouble is the dogma that to have a thought is to have an object before the mind. putnam and fodor (and many others) have distinguished two sorts of objects: those that are truly inner and thus 'before the mind' or 'grasped' by it, and those that identify the thought in the usual way. i agree that no objects can serve these two purposes, putnam (and some of the other philosophers i have mentioned) think the difficulty springs from the fact that an object partly identified in terms of external relations cannot be counted on to coincide with an object before the mind because the mind may be ignorant of the external relation, perhaps this is so, but it does not follow that we can find other objects which will insure the desired coincidence. for if the object isn 't connected with the world, we can never learn about the world by having that object before the mind; and for reciprocal reasons, it would be impossible to detect such a thought in another, so it seems that what is before the mind cannot include its outside connections-its semantics. on the other hand, if the object is connected with the world, then it cannot be fully before

the mind in the relevant sense. yet unless a *semantic* object can be before the mind *in its semantic aspect*, thought, conceived in terms of such objects, cannot escape the fate of sense data.

the basic difficulty is simple: if to have a thought is to have an object before the mind, and the identity of the object determines what the thought is, then it must always be possible to be mistaken about what one is thinking. for unless one knows *everything* about the object, there will always be senses in which one does not know what object it is. many attempts have been made to find a relation between a person and an object which will in all contexts hold if and only if the person can intuitively be said to know what object it is. but none of these attempts has succeeded, and i think the reason is clear. the only object that would satisfy the twin requirements of being before the mind and also such that it determines the content of a thought must, like hume's ideas and impressions, 'be what it seems and seem what it is'. there are no such objects, public or private, abstract or concrete.

the arguments of burge, putnam, dennett, fodor, stich, kaplan, evans, and many others to show that propositions can't *both* determine the contents of our thoughts *and* be subjectively assured are, in my opinion, so many variants on the simple and general argument i have just sketched. it is not just propositions that can't do the job. no objects could.

when we have freed ourselves from the assumption that thoughts must have mysterious objects, we can see how the fact that mental states as we commonly conceive them are identified in part by their natural history not only fails to touch the internal character of such states or to threaten first person authority; it also opens the way to an explanation of first person authority. the explanation comes with the realization that what a person's words mean depends in the most basic cases on the kinds of objects and events that have caused the person to hold the words to be applicable; similarly for what the person's thoughts are about. an interpreter of another's words and thoughts must depend on scattered information, fortunate training, and imaginative surmise, in coming to understand the other, the agent herself, however, is not in a position to wonder whether she is generally using her own words to apply to the right objects and events, since whatever she regularly does apply them to gives her words the meaning they have and her thoughts the contents they have. of course, in any particular case, she may be wrong in what

she believes about the world; what is impossible is that she should be wrong most of the time. the reason is apparent: unless there is a presumption that the speaker knows what she means, i.e. is getting her own language right, there would be nothing for an interpreter to interpret. to put the matter another way, nothing could count as someone regularly misapplying her own words. first person authority, the social character of language, and the external determinants of thought and meaning go naturally together, once we give up the myth of the subjective, the idea that thoughts require mental objects. this is an essay on an old topic, the relation between the human mind and the rest of nature, the subjective and the objective as we have come to think of them. this dualism, though in its way too obvious to question, carries with it in our tradition a large and not necessarily appropriate burden of associated ideas. some of these ideas are now coming under critical scrutiny, and the result promises to mark a sea change in contemporary philosophical thought. the present essay, while clearly tendentious, is not designed primarily to convert the skeptic; its chief aim is to describe, from one point of view, a fairly widely recognized development in recent thinking about the contents of the mind, and to suggest some of the consequences that i think follow from this development.

minds are many, nature is one. each of us has his own position in the world, and hence his own perspective on it. it is easy to slide from this truism to some confused notion of conceptual relativism. the former, harmless, relativism is just the familiar relativism of position in space and time. because each of us preempts a volume of space-time, two of us cannot be in exactly the same place at the same time. the relations among our positions are intelligible because we can locate each person in a single, common world, and a shared time frame.

conceptual relativism may seem similar, but the analogy is hard to carry out. for what is the common reference point, or system of coordinates, to which each scheme is relative? without a good answer to this question, the claim that each of us in some sense inhabits his own world loses its intelligibility.

for this reason and others, i have long held that there are limits to how much individual or social systems of thought can differ. if by conceptual relativism we mean the idea that conceptual schemes and

moral systems, or the languages associated with them, can differ massively—to the extent of being mutually unintelligible or incommensurable, or forever beyond rational resolve—then i reject conceptual relativism.<sup>1</sup> of course there are contrasts from epoch to epoch, from culture to culture, and person to person, of kinds we all recognize and struggle with; but these are contrasts which with sympathy and effort we can explain and understand. trouble comes when we try to embrace the idea that there might be more comprehensive differences, for this seems (absurdly) to ask us to take up a stance outside our own ways of thought.

we do not understand the idea of such a really foreign scheme. we know what states of mind are like, and how they are correctly identified; they are just those states whose contents can be discovered in well-known ways. if other people or creatures are in states not discoverable by these methods, it cannot be because our methods fail us, but because those states are not correctly called states of mind—they are not beliefs, desires, wishes, or intentions. the meaninglessness of the idea of a conceptual scheme forever beyond our grasp is due not to our inability to understand such a scheme, nor to our other human limitations; it is due simply to what we mean by a system of concepts.

many philosophers are not satisfied with arguments like these because they think there is another way in which conceptual relativism can be made intelligible. for it seems that we could make sense of such relativism provided we could find an element in the mind untouched by conceptual interpretation, then various schemes might be seen as relative to, and assigned the role of organizing, this common element. the common element is, of course, some version of kant's intuitions, hume's impressions and ideas, sense data, uninterpreted sensations, the sensuous given. kant thought only one scheme was possible; but once the dualism of scheme and content is made explicit, the possibility of alternative schemes is apparent. the idea that we have a choice of schemes is explicit in the work of c. i. lewis: 'there are, in our cognitive experience, two elements; the immediate data, such as those of sense, which are presented or given to the mind, and a form, construction, or interpretation, which represents the activity of thought.<sup>'2</sup>

<sup>&</sup>lt;sup>1</sup> i have argued for this in 'on the very idea of a conceptual scheme', essay 13 in *inquiries into truth and interpretation*.

<sup>2</sup> c. i. lewis, *mind and the world order*, 38. lewis declares that it is the task of philosophy 'to reveal those categorial criteria which the mind applies to what is given to it' (p. 36).

## the myth of the subjective 41

if we could conceive of the function of conceptual schemes in this way, relativism would appear to be an abstract possibility despite doubts about how an alien scheme might be deciphered: the idea would be that different schemes or languages constitute different ways in which what is given in experience may be organized. on this account, there would be no point of view from which we could survey such schemes, and perhaps no way we could in general compare or evaluate them; still, as long as we thought we understood the scheme-content dichotomy, we could imagine the unsullied stream of experience being variously reworked by various minds or cultures. in this way, it may be held, conceptual relativism can be provided with the element to which alternative schemes are related: that element is the uninterpreted given, the uncategorized contents of experience.

this picture of mind and its place in nature has defined many of the problems modern philosophy has thought it had to solve. among these problems are issues concerning knowledge: how we know about the external world, how we know about other minds, even how we know the contents of our own mind. but we should also include the problem of the nature of moral knowledge, the analysis of perception, and many troubling issues in the philosophy of psychology and the theory of meaning.

corresponding to this catalogue of problems is a long list of ways in which the supposed scheme-content contrast has been formulated. the scheme may be thought of as an ideology, a set of concepts suited to the task of organizing experience into objects, events, states, and complexes of such; or the scheme may be a language, perhaps with predicates and associated apparatus, interpreted to serve an ideology. the contents of the scheme may be objects of a special sort, such as sense data, percepts, impressions, sensations, or appearances; or the objects may dissolve into adverbial modifications of experience: we may be 'appeared to redly'. philosophers have shown ingenuity in finding ways of putting into words the contents of the given; there are those strange, verbless sentences like 'red here now', and the various formulations of protocol sentences over which the logical positivists quarreled.

expressing the content in words is not necessary and, according to some views, not possible. the scheme-content division can survive even in an environment that shuns the analytic-synthetic distinction, sense data, or the assumption that there can be thoughts or experiences

that are free of theory. if i am right, this is the environment provided by w. v. quine. according to quine's 'naturalized epistemology', we should ask no more from the philosophy of knowledge than an account of how, given the evidence we have to go on, we are able to form a satisfactory theory of the world. the account draws on the best theory we have: our present science. the evidence, on which the meanings of our sentences, and all our knowledge, ultimately depend, is provided by stimulations of our sense organs. it is these stimulations that provide a person with his only cues to 'what goes on around him'. quine is not, of course, a reductionist: 'we cannot strip away the conceptual trappings sentence by sentence'. nevertheless, there is according to quine a definite distinction to be made between the invariant content and the variant conceptual trappings, between 'report and invention, substance and style, cues and conceptualization'. for,

we can investigate the world, and man as a part of it, and thus find out what cues he could have of what goes on around him. subtracting his cues from his world view, we get man's net contribution as the difference. this difference marks the extent of man's conceptual sovereignty—the domain within which he can revise theory while saving the data/<sup>1</sup>

world view and cues, theory and data; these are the scheme and content of which i have been speaking.

what matters, then, is not whether we can describe the data in a neutral, theory-free idiom; what matters is that there should be an ultimate source of evidence the character of which can be wholly specified without reference to what it is evidence for. thus patterns of stimulation, like sense data, can be identified and described without reference to 'what goes on around us'. if our knowledge of the world derives entirely from evidence of this kind, then not only may our senses sometimes deceive us; it is possible that we are systematically and generally deceived.

it is easy to remember what prompts this view: it is thought necessary to insulate the ultimate sources of evidence from the outside world in order to guarantee the authority of the evidence for the subject. since we cannot be certain what the world outside the mind is like, the subjective can keep its virtue—its chastity, its certainty for us—only by being protected from contamination by the world.

<sup>&</sup>lt;sup>1</sup> this passage, and the quotations that precede it. are from w. v. quine, *word and object. 5.* 

the familiar trouble is, of course, that the disconnection creates a gap no reasoning or construction can plausibly bridge. once this starting point has been chosen, there is no saying what the evidence is evidence for, or so it seems. idealism, reductionist forms of empiricism, and skepticism loom.

the story is familiar; but let me continue in my breathless way through one more chapter. if the ultimate evidence for our schemes and theories, the raw material on which they are based, is subjective in the way i have described, then so is whatever is directly based on it: our beliefs, desires, intentions, and what we mean by our words. though these are the progeny of our 'view of the world'—indeed, taken together, they constitute our view of the world—nevertheless they too retain the cartesian independence from what they purport to be about that the evidence on which they are based had: like sensations, they could be just as they are, and the world be very different. our beliefs purport to represent something objective, but the character of their subjectivity prevents us from taking the first step in determining whether they correspond to what they pretend to represent.

instead of saying it is the scheme-content dichotomy that has dominated and defined the problems of modern philosophy, then, one could as well say it is how the dualism of the objective and the subjective has been conceived. for these dualisms have a common origin: a concept of the mind with its private states and objects.

i have reached the point i have been leading up to, for it seems to me that the most promising and interesting change that is occurring in philosophy today is that these dualisms are being questioned in new ways, or are being radically reworked. there is a good chance they will be abandoned, at least in their present form. what we are seeing is the emergence of a revised view of the relation of mind and the world.

the action has centered on the concept of subjectivity, what is 'in the mind'. let us start with what it is we know or grasp when we know the meaning of a word or sentence. it is a commonplace of the empirical tradition that we learn our first words (which at the start serve the function of sentences)—words like 'apple', 'man', 'dog', 'water'—through a conditioning of sounds or verbal behavior to appropriate bits of matter in the public domain. the conditioning works best with objects that interest the learner and are hard to miss by either teacher or pupil. but here is the point: this is not just a story

about how we learn to use words: it must also be an essential part of an adequate account of what words refer to, and what they mean.

needless to say the whole story cannot be this simple. on the other hand, it is hard to believe that this sort of direct interaction between language users and public events and objects is not a basicpart of the whole story, the part that, directly or indirectly, largely determines how the learner's words are related to things. one consequence is that the details of the mechanisms that constitute the causal chains from speaker to speaker, and spoken-of object to speaker to language learner, cannot in themselves matter to meaning and reference, the grasp of meanings is determined by the terminal elements in the conditioning process, and is tested only by the end product: use of words geared to appropriate objects and situations, this is perhaps best seen by noticing that two speakers who 'mean the same thing' by an expression need have no more in common than their dispositions to appropriate verbal behavior; the neural networks may be very different. the matter may be put the other way around: two speakers may be alike in 'relevant' physical respects, and yet they may mean different things by the same words because of differences in the external situations in which the words were learned. insofar, then, as the subjective or mental is thought of as supervenient on the physical characteristics of a person, and nothing more, meanings cannot be purely subjective or mental. as hilary putnam (misleadingly) put it, 'meanings ain't in the head'.<sup>4</sup> the point is that the correct interpretation of what a speaker means is not determined solely by what is in his head; it depends also on the natural history of what is in the head. putnam's argument hinges on rather elaborate thought experiments which some philosophers have found unconvincing. but the case can best be made by appealing directly to obvious facts about language learning and to facts about how we interpret words and languages with which we are unfamiliar.<sup>5</sup> the relevant facts have already been mentioned above; in the simplest and most basic cases, words and sentences derive their meaning from the objects and circumstances in whose presence they were learned. a sentence which one has been conditioned by the learning process to be caused to hold true by the presence of fires will (usually) be true when there is a fire present; a word one has been conditioned to hold

<sup>&</sup>lt;sup>4</sup> hilary putnam. 'the meaning of "meaning" ', 227.

<sup>&</sup>lt;sup>5</sup> i argue for this in essay 1.

applicable by the presence of snakes will refer to snakes. of course very many words and sentences are not learned this way; hut it is those that are that anchor language to the world.

if the meanings of sentences are propositions, and propositions are the objects of attitudes like belief, intention, and desire, then what has been said about meanings must hold true of all of the propositional attitudes. the point can be made without recourse to propositions or other supposed objects of the attitudes. for from the fact that speakers are in general capable of expressing their thoughts in language, it follows that to the extent that the subjectivity of meaning is in doubt, so is that of thought generally.

the fallout from these considerations for the theory of knowledge is revolutionary. if words and thoughts are, in the most basic cases, necessarily about the sorts of objects and events that commonly cause them, there is no room for cartesian doubts about the independent existence of such objects and events. doubts there can be, of course. but there need be nothing we are indubitably right about for it to be certain that we are mostly right about the nature of the world. sometimes skepticism seems to rest on a simple fallacy, the fallacy of reasoning from the fact that there is nothing we might not be wrong about to the conclusion that we might be wrong about everything, the second possibility is ruled out if we accept that our simplest sentences are given their meanings by the situations that generally cause us to hold them true or false, since to hold a sentence we understand to be true or false is to have a belief. continuing along this line, we see that general skepticism about the deliverances of the senses cannot even be formulated, since the senses and their deliverances play no central theoretical role in the account of belief, meaning, and knowledge if the contents of the mind depend on the causal relations, whatever they may be, between the attitudes and the world. this is not to deny the importance of the causal role of the senses in knowledge and the acquisition of language, of course.

it is an empirical accident that our ears, eyes, taste buds, and tactile and olfactory organs play an intermediate role in the formation of beliefs about the world. the causal connections between thought and objects and events in the world could have been established in entirely different ways without this making any philosophically significant difference to the contents or veridicality of perceptual belief. what is true is that certain beliefs caused by sensory experience are often veridical, and therefore often provide good reasons for further beliefs.

if this is right, epistemology (as apart, perhaps, from the study of perception, which is now seen to be only distantly related to epistemology) has no basic need for purely private, subjective 'objects of the mind', cither as uninterpreted sense data or experience on the one hand, or as fully interpreted propositions on the other. content and scheme, as remarked in the quotation from c. i. lewis, come as a pair; we can let them go together. once we take this step, no objects will be left with respect to which the problem of representation can be raised. beliefs are true or false, but they represent nothing. it is good to be rid of representations, and with them the correspondence theory of truth, for it is thinking there are representations that engenders intimations of relativism. (it is, of course, harmless to say true beliefs and thoughts are true because of the way the world is: they correctly 'represent' the world.)

there is an abundance of puzzles about sensation and perception; but these puzzles are not, as i said, foundational for epistemology. the question of what is directly experienced in sensation, and how this is related to judgements of perception, while as hard to answer as it ever was, can no longer be assumed to be a central question for the theory of knowledge. the reason has already been given: although sensation plays a crucial role in the causal process that connects beliefs with the world, it is a mistake to think it plays an epistemological role in determining the contents of those beliefs. in accepting this conclusion, we abandon the key dogma of traditional empiricism, what i have called the third dogma of empiricism. but that is to be expected: empiricism is the view that the subjective ('experience') is the foundation of objective empirical knowledge. i am suggesting that empirical knowledge has no epistemological foundation, and needs none.<sup>6</sup>

there is another familiar problem that is transformed when we recognize that beliefs, desires, and the other so-called propositional attitudes, are not subjective in the way we thought they were. the problem is how one person knows the mind of another. perhaps it is obvious that if the account i have sketched of our understanding of language, and its connection with the contents of thought, is correct, the accessibility of the minds of others is assured from the start. skepticism about the possibility of knowing other minds is thus ruled out. but to recognize this is not to answer the question

see essay 10.

what conceptual conditions we place on the pattern of thought that make it possible for an interpreter to progress from observed behavior to knowledge of the intentional attitudes of another person. that this question has an answer, however, is guaranteed by the fact that the nature of language and thought is such as to make them interpretable.<sup>7</sup>

it should not be assumed that if we cease to be bullied or beguiled by the scheme-content and subjective-objective dichotomies, all the problems of epistemology will evaporate. but the problems that seem salient will change. answering the global skeptic will no longer be a challenge, the search for epistemological foundations in preconceptual experience will be seen as pointless, and conceptual relativism will lose its appeal, but plenty of questions of equal or greater interest will remain, or be generated by the new stance. the demise of the subjective as previously conceived leaves us without foundations for knowledge, and relieves us of the need for them, but new problems then arise that cluster around the nature of error, for error is hard to identify and explain if the holism that goes with a nonfoundational approach is not somehow constrained. it is not problematic whether knowledge of the world and of other minds is possible, but it remains as a question how we attain such knowledge, and the conditions belief must satisfy to count as knowledge. these are not so much questions in traditional epistemology as they are questions about the nature of rationality, they are questions that, like the epistemological questions they replace, have no final answer; but unlike the questions they replace, they are worth trying to answer.

familiarity with many of the points i have been making is fairly widespread these days. but we have not fully appreciated the scope of the entailed revolution in our ways of thinking about philosophy. part of the reason for this failure may be traced to certain misunderstandings concerning the nature of the new antisubjectivism (as we may try calling it). here are three of the misunderstandings.

1. people have been persuaded of the dependence of meanings on factors outside the head by examples rather than by general arguments. there is therefore a strong tendency to suppose that the dependence is limited to the sorts of expressions that recur in the examples: proper names, natural kind words like 'water' and 'gold',

<sup>7</sup> see essay 1.

and indexicals.<sup>8</sup> but in fact the phenomenon is ubiquitous, since it is inseparable from the social character of language. it is not a local problem to be solved by some semantic trick; it is a perfectly general fact about the nature of thought and speech.

28. if mental states like belief, desire, intention, and meaning are not supervenient on the physical states of the agent alone, then, it has

been argued, theories that identify mental states and events with

physical states of and events in the body must be wrong. this is

suggested by putnam's claim that 'meanings ain't in the head', and it

is explicitly claimed by tyler burge and andrew woodfield.<sup>9</sup> the

argument assumes that if a state or event is identified (perhaps neces

sarily, if it is a mental state or event) by reference to things outside

of the body, then the state or event itself must be partly outside the

body, or at least not identical with any event in the body. this is

simply a mistake: one might as well argue that a sunburned patch of

skin is not located on the body of the person who is sunburned (since

the state of the skin has been identified by reference to the sun).

similarly, mental states are characterized in part by their relations to

events and objects outside of the person, but this does not show that

mental states are states of anything more than the person, nor that

they are not identical with physical states.

29. a third misunderstanding is closely related to the second. it is thought that if the correct determination of an agent's thoughts

depends, at least to some degree, on the causal history of those

thoughts, and the agent may be ignorant of that history, then the

agent may not know what he thinks (and, *mutatis mutandis*, what he

means, intends, etc.). the new antisubjectivism is thus seen as a

threat to first person authority—to the fact that people generally

know without recourse to inference from evidence, and so in a way

that others do not, what they themselves think, want, and intend. the

reaction has been to resort to maneuvers designed to insulate mental

states once more from their external determiners.

the maneuvers are not needed, for first person authority is not threatened. i may not know the difference between an echidna and a porcupine; as a result. i may call all echidnas i come across

<sup>&</sup>lt;sup>8</sup> tyler burge, 'individualism and the mental'; and hilary putnam. 'the meaning of "meaning" '.
' tyler burge, 'individualism and the mental', 111; and andrew woodfield, 'introduction', in woodfield (ed.), *thought and object*, p. viii.

porcupines. yet because of the environment in which i learned the word 'porcupine', my word 'porcupine' refers to porcupines and not to echidnas; this is what 1 think it refers to, and what i believe i see before me when i honestly affirm, 'that's a porcupine'. my ignorance of the circumstances that determine what i mean and think has no tendency to show that i don't know what i mean and think. to suppose otherwise is to show how strongly we are wed to the idea of subjective mental states that might be just as they are independent of the rest of the world and its history.

another reaction to the imagined threat to our real inner lives is to concede that beliefs and other mental states as we normally identify them are not truly subjective, but at the same time to hold that there are similar inner (or 'narrow') mental states that are, the idea might be, for example, that since nothing in my inner state or my behavior distinguishes between porcupines and echidnas, what i really believe when i see an echidna (or a porcupine) is that what is before me is an animal with certain general characteristics-characteristics that are in fact shared by porcupines and echidnas. the trouble is that since my word 'porcupine' refers only to porcupines, i apparently don't know what i mean when i say that's a porcupine'. this unattractive solution is unnecessary, for it is based on a confusion about what is 'inner'. since there is no present physical difference between my actual state and the state i would be in if i meant 'echidna or porcu pine' or 'animal with such and such properties' rather than 'porcu pine' and believed what i would then mean, it does not follow that there is no psychological difference. (there may be no relevant phys ical difference between the skin of someone who is sunburned and the skin of someone burned by a sunlamp, but there is a difference, since one state was and the other wasn't caused by the sun. psychological states are in this respect like sunburn.) so nothing stands in the way of saying i can know what i mean when i use the word 'porcupine' and what i believe when i have thoughts about porcupines, even though i can't tell an echidna from a porcupine. the psychological difference, which is just the difference between meaning and believing there is a porcupine before me, and meaning and believing there is a creature with certain common features of porcupines and echidnas, is exactly the difference needed to insure that i know what i mean and what i think. all that putnam and others have shown is that this difference does not have to be reflected in the physical state of the brain.

inventing a new set of truly 'inner', or 'narrow', psychological states is not, then, a way of restoring first person authority to the mental; quite the contrary. there remains the claim, however, that a systematic science of psychology requires states of the agent that can be identified without reference to their history or other connections with the outside world. otherwise, it is said, there would be no accounting for the fact that i, who may refer only to porcupines by my word 'porcupine', can no more tell the difference between a porcupine and an echidna than if (physically unchanged) i meant instead 'porcupine or echidna'.

the prospects for a scientific psychology are not directly relevant to the topic of this essay, and so we may disregard the question whether there are inner states of agents which might explain their behavior better than ordinary beliefs and desires. but it is relevant to consider whether there are states of mind which have a better claim to be called subjective than the propositional attitudes as these are usually conceived and identified.

two suggestions have been made. the more modest (to be found in the work of jerry fodor, for example) is that we might take as the true inner or solipsistic states selected states from among the usual attitudes, and modifications of these. thoughts about porcupines and echidnas would be eliminated, since the contents of such thoughts are identified by relations to the outside world; but admissible would be thoughts about animals satisfying certain general criteria (the very ones we use in deciding whether something is a porcupine, for example).<sup>10</sup>

such inner states, if they existed, would qualify as subjective by almost any standards: they could be identified and classified without reference to external objects and events; they could be called on to serve as the foundations of empirical knowledge; and the authority of the first person could plausibly apply to them. it seems clear, however, that there are no such states, at least if they can be expressed in words. the 'general features' or 'criteria' we use to identify porcupines are such as having four feet, a nose, eyes, and quills. but it is evident that the meanings of the words that refer to these features, and the contents of the concepts the words express, depend as much on the natural history of how the words and

<sup>&</sup>lt;sup>10</sup> jerry fodor, 'methodological solipsism considered as a research strategy in cognitive psychology'.

concepts were acquired as was the case for 'porcupine' and 'echidna'. there are no words, or concepts tied to words, that are not understood and interpreted, directly or indirectly, in terms of causal relations between people and the world (and, of course, the relations among words and other words, concepts and other concepts).

at this point one can imagine a proposal to the effect that there are inexpressible phenomenal criteria to which the publicly expressible criteria can be reduced; and here it is to be hoped that memories of past failures of such reductionistic fantasies will serve to suppress the thought that the proposal could be carried out. but even aside from nostalgic musings about phenomenalistic reduction, it is instructive to find the effort to make psychology scientific turning into a search for internal propositional states that can be detected and identified apart from relations to the rest of the world, much as earlier philosophers sought for something 'given in experience' which contained no necessary clue to what was going on outside. the motive is similar in the two cases: it is thought that a sound footing, whether for knowledge or for psychology, requires something inner in the sense of being nonrelational.

the second, and more revolutionary, suggestion was that the mental states needed for a scientific psychology, though roughly propositional in character, bear no direct relation to common beliefs, desires, and intentions." these states are, in effect, stipulated to be those that explain behavior, and they are therefore inner or subjective only in the sense of being characterized exclusively by what is physically beneath the skin. but there is no reason to suppose that people can tell without observation when they are in such states and so no reason to call them subjective.

in summary, i have made five connected points about the 'contents of the mind'.

*first,* states of mind like doubts, wishes, beliefs, and desires are identified in part by the social and historical context in which they are acquired; in this respect they are like other states that are identified by their causes, such as suffering from snow blindness or favism (a disease caused by contact with the fava bean).

second, this does not show that states of mind are not physical

<sup>&</sup>lt;sup>11</sup> this idea has been promoted by stephen stich, *from folk psychology to cognitive science*.

states of a person; how we describe and identify events and states has nothing directly to do with where those states and events are.

*third,* the fact that states of mind, including what is meant by a speaker, are identified by causal relations with external objects and events is essential to the possibility of communication, and it makes one mind accessible in principle to another; but this public and interactive aspect of the mind has no tendency to diminish the importance of first person authority.

*fourth*, the idea that there is a basic division between uninterpreted experience and an organizing conceptual scheme is a deep mistake, born of the essentially incoherent picture of the mind as a passive but critical spectator of an inner show. an adequate account of knowledge makes no appeal to such epistemological intermediaries as sense data, qualia, or raw feels. as a result, global skepticism of the senses is not a position we can coherently entertain.

*finally*, i have argued against the postulation of 'objects of thought', whether these are conceived on the model of sense data, or as propositional in character. there are the many states of mind, but their description does not require that there be ghostly entities that the mind contemplates. to dispense with such entities is to eliminate rather than solve a number of vexing problems. for we cannot ask how such objects can represent the world if there are no such objects, nor can we be puzzled by the question how the mind can be directly acquainted with them.

what remains of the concept of subjectivity? so far as i can see, two features of the subjective as classically conceived remain in place. thoughts are private, in the obvious but important sense in which property can be private, that is, belong to one person. and knowledge of thoughts is asymmetrical, in that the person who has a thought generally knows he has it in a way in which others cannot. but this is all there is to the subjective. so far from constituting a preserve so insulated that it is a problem how it can yield knowledge of an outside world, or be known to others, thought is necessarily part of a common public world. not only can others often learn what we think by noting the causal dependencies that give our thoughts their content, but the very possibility of thought demands shared standards of truth and objectivity.

# what is present to the mind?

there is a sense in which anything we think about is, while we are thinking about it, before the mind. but there is another sense, well known to philosophers, in which only some of the things we can think about are said to be before the mind. some are before the mind when we think about anything at all; others we do not have to think about for them to be before the mind. both of these are things supposed to be so directly before the mind that it is impossible to misidentify them, or we can misidentify them only if we do not know what we think; in this they differ from ordinary physical objects or indeed any other objects at all, all of which are easy to misidentify. we cannot mistake these entities for others simply because it is these objects that fix the contents of our thoughts. it is things of this second sort i want to discuss.

some things of this second sort are the objects, so called, of desires, beliefs, intentions, worries, and hopes; they are the propositions to which we have the various attitudes, the thoughts (as frege named them) which our sentences express. we must also include the constituents of propositions, such things as properties and relations.

the objects i have mentioned have a special relation to the mind: it is only through the mind that we can know them. since they are abstract, they cannot be accessed by the senses. they have no causal powers, and so cannot act on, or be acted on by, our minds, our brains, or us.

this raises the question what sort of relations we or our minds are thought to have to these objects when we say they are before the mind. we have many words to express our relations to propositions: we grasp them when we understand a sentence, we entertain them, reject them, wish they were true, hope they aren't, or intend to make them true. but what sort of relations are these? they sound like psychological relations, as if there were some sort of mental

transaction between us and these entities. but what kind of business can be transacted with an abstract object?

the propositional objects of the mind and their constituents are supposed, then, to have these two properties: they *identify a* thought by fixing its content; and they *constitute* an essential aspect of the psychology of the thought by being grasped or otherwise known by the person with the thought. it is a problem how these two properties can be reconciled. my thesis is that they cannot.

here is the main difficulty. i take for granted that for the most part we do know what we think, even though there are departures from total self-knowledge. (i use 'think' to cover all the propositional attitudes.) but if a thought is constituted the thought it is by the mind's knowledge of the identifying object, then someone knows what thought she is thinking only if she knows which object she has in mind. yet there seems to be no clear meaning to the idea of knowing which object one has in mind. the trouble is that ignorance of even one property of an object can, under appropriate circumstances, count as not knowing which object it is. this is the reason philosophers who have wanted to found knowledge on infallible identification of objects have sought objects that, like hume's impressions and ideas, 'are what they seem and seem what they are'-that is, have all and only the properties we think they have, alas, there are no such objects, every object has an infinity of logically independent properties, even those objects, like numbers, all of whose 'essential' properties we specify.<sup>1</sup>

<sup>1</sup> in *the interpretation of frege's philosophy*, 50, michael dummett describes this feature of fregean senses: 'a sense cannot have any features not discernible by reflection on or deduction from what is involved in expressing or in grasping it. only that belongs to the sense of an expression which is relevant to the determination of the truth value of the sentence in which it occurs; if we fail to grasp some features of its contribution to the truth-conditions of certain sentences, then we fail fully to grasp its sense, while, on the other hand, any aspect of its meaning that does not bear on the truth-conditions of the sentences containing it is no part of its sense. it cannot be, therefore, that the sense has all sorts of other features of senses we cannot fail to detect to their 'internal properties', but it is not clear on what non-question-begging principle such properties are to be told from others.

steve yablo asks (in a private communication) why we couldn't pay attention to some only of the properties of the objects, perhaps the 'essential' properties. no, not if we are to stick to the idea that the contents of the thought are individuated by the object. for if two objects have the same 'essential' properties, but differ in some other property, they are different objects, and so the thoughts must differ.

so perhaps it isn't the *object* that is before the mind, but only some *aspect* of it. we might get the object wrong, but it wouldn't matter. but then the aspect is the object of thought, not the object of which it is the aspect; and the same difficulties arise anew.

#### what is present to the mind 55

recent discussion of de re belief brings out the point. if we agree with russell that a person cannot form a judgement about an object unless he knows which object it is, or (to put it in another way that russell favored) it is an object with which the person is acquainted, and that this demands, in the case of propositions, that this special relation hold between the judge and each part of the proposition judged, then there is a problem about attitudes like the following: there is a recipe for making corn bread that joan believes is easy. the truth of this attribution not only demands that there be a recipe for making corn bread, but it also seems to require that joan know which recipe it is—or be somehow acquainted with it. when guine first emphasized the distinction between *de re* and *de dicto* belief sentences in modern terms, he was inclined to see *de re* reference as an island of clarity in the opaque sea of intensionality.<sup>2</sup> subsequent developments led to a change of mind. in 'intensions revisited' he wrote, the notion of knowing or believing who or what someone or something is, is utterly dependent on context.<sup>3</sup> what led to the change of mind was the difficulty in explaining the relation between a person and an object that would justify the claim that the person knew which object it was. a number of attempts have been made to clarify the relation: f011esdal declared only 'genuine' names could enter into it; kripke spoke of 'rigid designators'; david kaplan called the elect names 'vivid'.4 gareth evans studied the problem in depth; he thought, with others, that the only psychological relation that could count as providing the requisite sort of 'fundamental identification' of an object was demonstrative identification. in such a case alone could one say that the object was part of the proposition entertained. following russell, evans concluded that when a person thinks he is entertaining a singular thought, but is using a nonreferring name, there is no proposition for him to contemplate, and therefore no thought that he has. if he uses a sentence containing a nonreferring name, he expresses no thought.5 if, like me, you have trouble feeling confidence in the criteria for genuine cases of 'fundamental identification', you will appreciate why russell limited such cases to situations in which the mind is directly acquainted with its

<sup>&</sup>lt;sup>2</sup> w. v. quine, 'quantifiers and propositional attitudes'. 3 w.

v. quine, 'intensions revisited'. 121.

<sup>&</sup>lt;sup>4</sup> dagfinn f0llesdal, 'knowledge. identity, and existence': saul kripke. 'naming and necessity'; david kaplan, 'quantifying in'.

<sup>&</sup>lt;sup>5</sup> gareth evans, the varieties of reference.

objects, something he thought was possible only with sense data (and perhaps with oneself).

what lies behind some of these attempts to characterize the special relation between the mind and its objects is, of course, the cartesian drive to identify a sort of knowledge which is guaranteed against failure. if this search is combined with the assumption that all knowledge consists in the mind being in psychological contact with an object, then objects must be found about which error is impossible—objects that must be what they seem and seem what they are.

there simply are no such objects. not even appearances are everything we think they are. nor can the 'aspects' of sense data, if they really are objects, be protected from one or another sort of misidentification. the reason quine is right in thinking we cannot pick out 'the' relation that constitutes knowing which object some object is is simply that *any* property of an object may, under suitable conditions, be considered the relevant identifier.

i have dwelt briefly on the problem as it has been studied in relation to proper names because it is there that philosophers seem to have come closest to appreciating the nature of the difficulty. but if my diagnosis is correct, the problem really has nothing special to do with proper names; it is a perfectly general problem about the objects of the mind. if the mind can think only by getting into the right relation to some object which it can for certain distinguish from all others, then thought is impossible. if a mind can know what it thinks only by flawlessly identifying the objects before it, then we must very often not know what we think.

why, if they make so much trouble, do we suppose there *are* propositional objects of thought? well, for one thing we certainly talk as if there were such entities: we think deep thoughts, share thoughts, discard and embrace beliefs, entertain, consider, reflect on, contemplate ideas and propositions—one could go on for a long time. these are the sort of remarks we have been taught, with reason, to view with ontological suspicion. but it is much harder to take lightly the problem we get into when we attempt to give a systematic account of what certain sentences mean—the ones we use to attribute thoughts to thinkers. for it is difficult to see how ordinary sentences like 'paul believes that the koh-i-noor diamond is one of the crown jewels' can be analyzed except as relating paul to some entity picked out by the phrase 'that the koh-i-noor diamond is one of the crown jewels', or perhaps by the sentence 'the koh-i-noor diamond is one

of the crown jewels'. plenty of attempts have been made to avoid taking 'believes' as a relational term, that is, as relating two objects, a believer and something else, but none of these attempts is, as far as i know, successful.

one proposal is to treat the rest of the sentence after 'believes' as a complex adverb. the reading might be something like: 'paul believes in a that-the-koh-i- noor-diamond-is-one-of-the-crownjewels fashion'. but this is an unsatisfactory suggestion, since no one has any idea how to derive the meaning of such an adverbial modifier from the meanings of the constituent words. yet it is intuitively obvious that we understand the sentences that follow the 'that' in belief sentences because we understand the constituent words, if the meanings of such contained sentences (which give the 'contents' of propositional attitudes) are not constructed from the meanings of their parts, they must have to be learned independently—as if they were new, often very long, words. it seems clear that this is a wrong idea, and probably an impossible one, since any declarative sentence can feature as a content-sentence, and there is an unlimited, and so presumably unlearnable, totality of these.

other suggestions along similar lines have been made fairly often, but no one has ever shown how to implement such suggestions by incorporating them in a developed semantic theory.<sup>6</sup> the governing principle seems clear: apart from sentences, whenever we recognize a grammatical category to which we must assign an infinity of expressions, ontology is required. there must be an infinity of objects that can be referred to either indexically or by means of descriptions; these are the two devices available to enable a finite vocabulary to cope with any of a potential infinity of objects. objects allow us to manage adjectives; events do the same for at least some adverbs. numbers do the job when we want to measure.<sup>7</sup>

there is, then, no plausible alternative to taking belief sentences as relational, and therefore no alternative to taking the content sentence ('the koh-i-noor diamond is one of the crown jewels' plus, perhaps, the word 'that' when it is present) as a singular term

<sup>&</sup>lt;sup>6</sup> for a discussion of a number of attempts to eliminate the ontological commitment to 'objects of thought' in the analysis of belief sentences, and other sentences that attribute propositional attitudes, see essay 7 in *inquiries into truth and interpretation*.

<sup>7</sup> see w. v. quine, 'events and reification'.
which, by referring to an appropriate entity, specifies the relevant belief. $^{8}$ 

apparently we have a dilemma. on the one hand, there is the fact that to have a belief or other propositional attitude is to be related to an object of some sort; on the other hand, there is the fact that there seems to be no satisfactory account of the psychological relation a person must be in to the appropriate object in order to have the attitude. the difficulty in giving such an account hinges on the idea that since a person generally knows what he thinks, he must be directly acquainted with, or be able in some special way to identify or individuate, the object or objects that define (give the contents of) his thought.

if we rid ourselves of preconceptions, i think it is easy to see where we have gone wrong. it does not follow, from the facts that a thinker knows what he thinks and that what he thinks can be fixed by relating him to a certain object, that the thinker is acquainted with, or indeed knows anything at all about, the object. someone who attributes a thought to another must, i have argued, relate that person to some object, and so the attributer must, of course, identify an appropriate object, either by pointing to it or by describing it. but there is no reason why the attributer must stand in any special relation to the identifying object; all he has to do is refer to it in the way he refers to anything else.

we specify the subjective state of the thinker by relating him to an object, but there is no reason to say that this object itself has a subjective status, that it is 'known' by the thinker, or is 'before the mind' of the thinker. this consequence was already implicit in some analyses of sentences about propositional attitudes, for example the suggestion once entertained by carnap, and discussed by quine, that belief sentences be taken to relate a believer to a sentence of the attributer. thus quine, who holds that a cat can have a belief, points out that there is no reason to suppose a cat is acquainted with the sentence 'food's on'just because it can correctly be said that the cat believes that food's on. for those who doubt that cats have propositional attitudes, the

<sup>&</sup>lt;sup>s</sup> stephen schiffer rightly emphasizes that the possibility of 'compositional semantics' (the idea that the meanings of complex expressions must he seen as a function of the mean ings of the expressions of which they are composed) depends crucially on the relational analysis of sentences that are used to attribute attitudes. he despairs of finding a satisfac tory analysis for reasons i do not share, and so abandons hope for a compositional seman tics. stephen schiffer. *remnants of meaning*.

same point can be made by remarking that we may identify a belief of sebastian's by saying he believes that naples is north of san francisco, though he doesn't know a word of english. my point here isn't that belief sentences relate believers to sentences, but that this familiar proposal assumes that the objects used to identify a belief may not be within the ken of the believer. once we grant this possibility, we are free to divorce the semantic need for content-specifying objects from the idea that there must be any objects at all with which someone who has an attitude is in psychic touch.

here is an analogy: consider weight. some things weigh more than others; some things weigh nothing; occasionally two things weigh the same. one thing may weigh twice what another does. these relations among objects are what we wish to report when we assign weights to them. introducing a standard does not alter the situation. thus a monetary pound in the time of william i weighed the same as 12 ounces; it took 20 pennies to weigh an ounce; and a penny weighed the same as 32 grains of wheat taken from the middle of the ear. thus a pound weighed the same as 7,680 grains of wheat. all that we wish to say about how much things weigh can be put in terms of these comparisons: for example, the koh-i-noor diamond in its present condition weighs about the same as  $15'_{4}$  pennies, or 490 grains of wheat.

all these comparisons can be tedious, and in any case we make the relevant comparisons perspicuously by using numbers. so the convenient thing is to settle on a way of representing the relations among objects with respect to weight by using numbers directly. thus we say the koh-i-noor diamond weighs 109 carats or 345 grams. but talk of this sort does not require us to include carats or grams in our ontology. the only objects we need are the numbers and the things that have weight. to say the weight in carats of the koh-inoor diamond is 109 does not commit us to *weights* as objects: it is just to assign the number 109 to the diamond as a way of relating it to other objects on the carat scale.

seen this way, talk of how much things weigh is relational: it relates objects to numbers, and so to one another. but no one supposes the numbers are in any sense intrinsic to the objects that have weight, or are somehow 'part' of them. what are basic are certain *relations* among objects: we conveniently keep track of these relations by assigning numbers to the objects, and remembering how the relations among the objects are reflected in the numbers.

one important aspect of numerical measurement is that typically

some only of the properties of numbers are relevant to their use in reporting relative weights. thus it *is* relevant that whatever numbers we use to keep track of weights preserve ratios: if one thing weighs twice as many grams as another, it must also weigh twice as many pounds. on the other hand, the absolute size of the number is irrelevant: measurement in pounds, grains of wheat, and grams yields different numbers, but the same *relative* weights when compared to other objects.<sup>9</sup>

the analogy with beliefs is this. just as in measuring weight we need a collection of entities which have a structure in which we can reflect the relations between weighty objects, so in attributing states of belief (and other propositional attitudes) we need a collection of entities related in ways that will allow us to keep track of the relevant properties of and relations among the various psychological states.

in thinking and talking of the weights of physical objects we do not need to suppose there are such things as weights for objects to have. similarly, in thinking and talking about the beliefs of people we needn't suppose there are such entities as beliefs. nor do we have to invent objects to serve as the 'objects of belief or what is before the mind, or in the brain. such invention is unnecessary because the entities we mention to help specify a state of mind do not have to play any *psychological* or epistemological role for the person in that state, just as numbers play no physical role. as a consequence, there is no reason to conclude, from the thinker's lack of knowledge of the entities we use to track his thoughts, that *he* may not know what he thinks.

the suggestion i am proposing about the nature of the propositional attitudes applies directly to a problem that has troubled a number of philosophers in recent years. there are convincing arguments to show that the correct determination of the contents of beliefs (and meanings and other propositional attitudes) depends in part on causal connections between the believer and events and objects in the world of which he may be ignorant. a now familiar example is putnam's twin earth case.<sup>10</sup> we are invited to imagine

<sup>&</sup>lt;sup>9</sup> a comparison between how we attribute beliefs and numerical measurement has been made by paul churchland in *scientific realism and the plasticity of mind*, 105. but he makes the mistake of supposing we can take a phrase like '345 grams' in 'the koh-i-noor diamond weighs 345 grams' as an *adverb* of "weighs", and thus get rid of the onto-logical problem (he would say the same about the sentence that follows 'believes'). but as we have seen, this suggestion cannot be supported by a serious semantics.

<sup>&</sup>lt;sup>10</sup> hilary putnam, the meaning of "meaning"'.

#### what is present to the mind 61

that there is a twin to our earth which is, in all immediately discernible respects, identical with our earth. on twearth is my doppelganger, molecule for molecule the same (that is, similar in 'relevant respects'), having been exposed to the same conditioning, and having exactly the same linguistic dispositions, as i have. yet one of us believes it is water he sees before him (me) when the other believes it is twater. the explanation is that where there is water on earth, there is twater on twearth, though no one has yet detected the difference, putnam argues that there is no inner, or psychological, difference between me and my twin. neither of us has any reason to say he believes one thing rather than the other. therefore (putnam continues) neither of us knows what he believes. (i don't know whether i believe this is water or twater.) so there may be, and perhaps always are, nonsubjective factors, factors unknown to the thinker, which decide what the 'object of thought' is. if the identity of the 'object of thought' is partly dependent on factors of which the person who has the thought is ignorant, doesn't it follow that the person may not know what he thinks?

the answer is that it doesn't follow. it would follow if the object used to identify my thought were something i had to be able to discriminate in order to know what i think. but this is the assumption i have urged we should abandon. what i see before me i believe to be water; i am in no danger of thinking it is twater, since i do not know what twater is. if i am on earth, i also believe i think i am seeing water, and in this i am right, if i were without my knowledge transferred to twearth i would believe twater was water-a mistake. in both cases i would know what i believed. of course my twin on twearth refers to twater by the word 'water'. he refers to twater by using the word 'water' because he thinks it is twater, this is also what he thinks he thinks, so neither of us would be wrong about what he thought, the possibility of error, or of failure to distinguish one's own state of mind, due solely to the external elements that help determine that state of mind, is intelligible only on the supposition that having a thought requires a special psychological relation to the object used to identify the state of mind.11

the point of this exercise may surprise us. it is that subjective states are not supervenient on the state of the brain or nervous

<sup>&</sup>lt;sup>11</sup> for further discussion of this point, see essay 2.

system: two people may be in similar physical states<sup>12</sup> and yet be in dissimilar psychological states. this does not mean, of course, that mental states are not supervenient on physical states, for there must be a physical difference *somewhere* if psychological states are different. but the interesting physical difference may not be in the person; like the difference between water and (water, it may be (we are supposing) elsewhere.

many philosophers faced with this result have decided that beliefs and other so-called propositional attitudes, as we usually identify them, are not quite as subjective as they were thought to be. thus jerry fodor, in recommending 'methodological solipsism' to the psychologist, thinks he is recommending that the psychologist deal with *truly* subjective states, states whose identities are determined only by what is in the head. david kaplan and daniel dennett have made similar suggestions.<sup>13</sup> stephen stich's idea is along the same lines, but more portentous, since he thinks we must give up folk psychology entirely if we want to have a serious science of human behavior.<sup>14</sup> i agree that the concepts of belief, desire, intention, and the like are not suited to a science like the physical sciences, and i agree that one of the reasons is that mental states are partly identified on the basis of their causes and effects, but since the same is true of human actions as we normally identify them, there seems to be no chance of more 'scientific' explanations in this area (assuming that what we mean by a 'scientific' explanation is the kind of explanation we hope for from physics).

but whatever we think about the scientific future of folk psychology, there is no good reason for claiming that beliefs and the other propositional attitudes are not truly subjective states. we have discovered no reason for saying that thinkers don't generally know what they think, or that there isn't always a presumption in favor of their being right when they disagree with others about what is on their minds. and this seems to be as good a test of subjectivity as we have.

there are two important points i have left hanging, and i shall devote a few remarks to them. first, i have not said anything about

<sup>14</sup> stephen stich, from folk psychology to cognitive science.

<sup>12</sup> not. of course, *identical* physical states, as is often erroneously said. for physical states cannot in the relevant sense be identical unless they inhere in the same object.

<sup>&</sup>lt;sup>13</sup> jerry fodor, 'methodological solipsism considered as a research strategy in cognitive psychology', david kaplan, 'quantifying in', daniel dennett, 'beyond belief.

#### what is present to the mind 63

what the objects are to which we refer when we want to specify what someone believes. and, second, i have argued that the discovery that what is believed is partly settled by facts of which the believer may be ignorant does not show he doesn't know what he believes. but this is only a negative point; it does not even hint at *why* there is a presumption—usually correct—that he does know. these two points are closely related.

first, then, what are the objects we name or describe in order to characterize states of mind? well, what do we *say*? we say things like 'paul believes the koh-i-noor diamond is one of the crown jewels'. the words 'believes that the koh-i-noor diamond is one of the crown jewels' characterize an aspect of paul's state of mind. the relational word is 'believes', and what follows names the object, not of thought, i have insisted, but the object which in some regular way indicates paul's state of mind. i have in the past suggested that we take the word 'that' in such sentences as a demonstrative that picks out or refers to the next utterance of the speaker who is doing the attributing. the following gives the idea: paul believes what i would believe if i were sincerely to assert what i say next.

the koh-i-noor diamond is one of the crown jewels.

since i am not trying to make subtle semantic points, you may take as the 'object' either my actual utterance, or the sentence of which it is an utterance (relativized to a time and a speaker). if you wish, you even may take the object to be a proposition. since utterances, sentences, and propositions are so closely related, the chances are if one choice will serve, the others can be made to serve. but utterances have certain prima facie advantages, since they are non-abstract, and so come with a speaker, a time, and a context attached. so i will assume we have settled on utterances, the very utterances that are produced in attributing attitudes, as the objects that serve to individuate and identify the various states of mind.

such objects serve very well in their role of characterizing states of mind. there certainly are as many different utterances (or potential utterances) at our disposal as there are states of mind we are able to distinguish in attributing them—more, in fact. even the special difficulties having to do with demonstrative or indexical expressions are not insoluble if we give up the idea that we are trying to identify objects with which those who have attitudes are in mysterious psychic touch.

utterances are related to each other in much the same way beliefs are: by relations of entailment and evidential support. utterances, like beliefs, are true and false. aside from complications due to indexical elements, we identify a belief by uttering a sentence that has the same truth conditions as the belief it is used to identify. nor is any of this surprising since we often express our beliefs by uttering sentences with the truth conditions of the belief we are expressing.

at this point you may think i am coming dangerously close to restoring the very theory that made all the trouble in the first place. for why not go on to say that since utterances have determinate meanings, and it is meanings that match up with belief states, the objects we are naming when we utter content sentences in the context of attributing beliefs just are the meanings of those content sentences, that is, propositions. this would fix it so that when a frenchman attributed the same state of mind to paul as i did, we would both be naming the same object: this would not be the case on the theory i was just considering, for the frenchman's relevant sentence was not mine.

i have only marginal objections in the present context to taking this step, for it is not the step that makes the trouble. (it may make some other kind of trouble, meanings and propositions being the tricky things they are.) the trouble we have been concerned with here sprang from the identification of the object used to characterize a state of mind with an object that the mind 'knows' or is 'acquainted with', an 'object of thought'. if we avoid this identification, we can stay out of the difficulties i have been exploring, but if we avoid this identification, neither do we gain anything by the steps from utterances to sentences to meanings or propositions. so as a help in keeping my main point in mind, it is well to stick to utterances. there will then be no danger of supposing that in general the believer is acquainted with the objects used to characterize his states of mind. it should not bother us that the frenchman and i use different utterances to characterize the same state of paul's mind: this is like measuring weight in carats or ounces: different sets of numbers do the same work.

this last point directs our attention to a larger issue. when we use numbers to keep track of the relations among weights and lengths and temperatures, we are not apt to respond to the fact that different sets of numbers do as well as others in keeping track of all that is relevant empirically by complaining that weights or lengths or temperatures

#### what is present to the mind 65

are not 'real'. we know there is no contradiction between saying that the temperature of the air is  $32^{\circ}$  fahrenheit and saying that it is  $0^{\circ}$ celsius; there is nothing in this 'relativism' to show that the properties being measured are not real. curiously, though, this conclusion has repeatedly been drawn. john searle, for example, finds it incomprehensible that either of two quite different interpretations might correctly be put on the same thought (or utterance) of a person.<sup>15</sup> yet in the light of the considerations put forward here, this amounts merely to the recognition that more than a single set of one person's utterances might be equally successful in capturing the contents of someone else's thoughts or speech, just as numbers can capture all the empirically significant relations among weights or temperatures in infinitely many different ways, so one person's utterances can capture all the significant features of another person's thoughts and speech in different ways. this fact does not challenge the reality of the attitudes or meanings thus variously reported. jerry fodor is another philosopher who thinks that holism, or the indeterminacy of translation that is associated with it, threatens realism with respect to the propositional attitudes.<sup>16</sup> this is the same mistake. indeterminacy of translation means that different sets of utterances (or sentences, as quine has it) do equally well in interpreting a speaker's language (or thoughts); this does not suggest that the states of mind of the speaker or thinker thus captured are somehow vague or unreal. holism maintains that the contents of speech and thought depend on the relations among meanings and among thoughts. but again there is nothing in this claim to threaten the reality of the states that are related in these ways. the threat to the reality of thoughts and meanings that searle and fodor think they detect is in fact a quite different threat, a threat to the assumption that the entities used to identify thoughts and meanings are somehow 'grasped' by the mind, and so, if the entities are different, the thoughts themselves must be different. it is as if the 'difference' between being a yard long and 36 inches long in a yardstick were a difference in the yardstick itself.

i come now to the last question: given that a correct view of the way in which the 'objects of thought' determine the identity of the various states of mind does not threaten first person authority, what accounts for that authority?

<sup>&</sup>lt;sup>15</sup> john searle. 'indeterminacy and the first person'.

<sup>&</sup>lt;sup>16</sup> jerry fodor. *psychosemantics*.

one difficulty was created by grammar and false inferences from it: these led to the idea of an inner object known in a privileged way. this obstacle was removed by showing that there was no reason to suppose there are such objects. a second difficulty sprang from the conviction that the truly subjective—that of which the person has privileged knowledge—can owe none of its subjective quality to connections with the outside world. i have argued that though an interpreter must, if he is to get things right, look to relations between the mind he is interpreting and its environment, this does not prejudice the self-knowledge of the knower.

having removed the obstacles, there is little left to say about how we know what we think. in the interesting, and originally puzzling cases, there is no way we know-for there is no evidence to be sought, no inner object to be scrutinized, no competing hypothesis to be weighed. the point comes out if we consider this situation: suppose i say, 'i believe the koh-i-noor diamond is a crown jewel'. or don't just suppose it, for i do say this, and suppose, as is the case, that i know what the words i have just uttered mean, and that i am making a sincere assertion. finally, let us suppose that both you and i agree on these points, namely that i did utter those words and that in doing so i was sincerely uttering words i understood. from these suppositions it follows that i know what i believe, but it does not follow that you know what i believe, the reason is simple; you may not know what i mean. your knowledge of what my words mean has to be based on evidence and inference: you probably assume you have it right, and you probably do. nevertheless, it is a hypothesis. of course, i may not know what i mean by those words either. but there is a *presumption* that i do, since it does not make sense to suppose i am generally mistaken about what my words mean; the presumption that i am not generally mistaken about what i mean is essential to my having a language-to my being interpretable at all. to appeal to a familiar, though often misunderstood, point: i can do no better, in stating the truth conditions for my utterance of the sentence the koh-i-noor diamond is a crown jewel' than to say it is true if and only if the koh-i-noor diamond is a crown jewel. if i say this, i utter a tautology, but if you give the truth conditions of my utterance in the same words, you are making an empirical claim, though probably a true one.17

<sup>&</sup>lt;sup>17</sup> this paragraph summarizes the conclusion of essay 1 in this volume.

#### what ix present to the mind 67

according to dummett, brentano 'refused to admit that a mental act. . . had any inner object distinct from the external one, namely a mental representation . . . by which the external object was presented to the mind'.<sup>18</sup> dummett points out that this leaves brentano with the problem of thoughts (or apparent thoughts) about objects that don't exist, a problem, dummett remarks, brentano 'did not succeed in resolving'. but the problem is easily solved if we give up the idea that there are inner objects or mental representations in the required sense. there is no need to suppose that if there are no such inner objects, only outer objects remain to help us identify the various states of mind. the simple fact is that we have the resources needed to identify states of mind, even if those states of mind are, as we like to say, directed to nonexistent objects, for we can do this without supposing there are any objects whatever before the mind.

<sup>&</sup>lt;sup>18</sup> michael dummett, *the interpretation of frege's philosophy*, 57. dummett refers us to franz brentano's *psychology from an empirical standpoint*, 79.

## 5 indeterminism and antirealism

antirealism is a manifestation of the irrepressible urge in western philosophy to insure that whatever is real can be known: antirealism attempts to achieve this by reading out of existence whatever it decrees lies beyond the scope of human knowledge, so parmenides should be counted an early and extreme antirealist, since he declared that the globular, homogeneous, changeless one was all that was real on the ground that it was the only possible object of knowledge. plato must, of course, be counted an antirealist, since he held that the physical world, and all in it, is less than real because it cannot be known. most reductive isms should count as forms of antirealism: idealism, pragmatism, empiricism, materialism, behaviorism, verifi-cationism. each tries to trim reality down to fit within its epistemol-ogy. each of these positions offers consolations: ordinary objects like tables and dinnerware are real, we are told, but exist only in the mind; physical objects are nothing but permanent possibilities of sensation; mental states are nothing but patterns of behavior; intentional phenomena are nothing but physical events and objects; etc. we are allowed the terminology of our old ontology as long as we agree to accept only what can be cobbled together out of entities or experiences we can for certain know. but these sops to skepticism should not deceive us; antirealisms remain sour grape philosophies. their motto is: if you can't grasp the grapes (in some approved sense), they aren't just sour; they were never there in the first place. for reasons i shall mention presently, some antirealisms are better expressed in terms of epistemic limitations on the concept of truth. thus it may be held that when our epistemic faculties are deficient with respect to determining the truth or falsity of some sentence, we should rule that the sentence has no truth value, or that we should

employ some reduced sense of truth. the outcome is the same: the real or the true is cut down to the size of a favored form of knowledge.

i do not see realism as the alternative to antirealism. we can reject one or another version of antirealism on the grounds that the arguments for it fail without having to endorse some vague position called realism. one common characterization of realism is this: there is something in or about the world that makes our utterances or assertions or thoughts true when they are true, whether or not we have the power to determine their truth. most criticisms of this formulation fasten, understandably, on the difficulty in spelling out what 'power' and 'determine' mean in this context, but i don't think we need to get in this deep, since i do not think we understand the prior claim that there is something in or about the world that makes our thoughts or assertions true when they are true. the fact is that no one has ever been able to say in a nontrivial way what sort of 'thing' it is that makes a sentence (or other truth bearer) true. some sentences or utterances or beliefs are true, and there are a lot of things in the world, but it explains nothing to say that the things 'make' the truth vehicles true. the result is that realists are left holding the concept of truth, but can't explain how reality accounts for it. the failure of correspondence theories to give substance to the concept of truth at the same time demonstrates the vacuity of characterizations of realism in terms of correspondence.

here i discuss those versions of antirealism that throw doubt on the reality of mental states and events in the light of the indeterminacy of translation or interpretation. i am equally concerned with those who argue that *if we* accept the indeterminacy of interpretation, then we must be in doubt about the status of propositional attitudes or at least the status of attributions of propositional attitudes. since i accept the thesis of the indeterminacy of interpretation, i would be aggrieved to find that it entails some sort of antirealism. but i do not think it does, as i shall try to explain. i think neither that indeterminism shows that propositional attitudes are less than fully real (whatever that may mean) nor that we must modify the concept of truth when we come to talk of propositional attitudes. in other words, many of our beliefs and statements about what people believe, intend, desire, and hope for are true, and they are true because people have those attitudes.

three considerations seem to stand in the way of accepting this thesis. the first is a form of scientism. the propositional attitudes do

#### indeterminism and antirealism 71

not seem suited to incorporation into a unified scientific view of the world, thus quine has said that 'the essentially dramatic idiom of propositional attitudes' has no place in serious science,<sup>1</sup> or that our mentalistic vocabulary is 'stubbornly at variance with scientific patterns'. i think quine is right in holding that the mentalistic vocabulary cannot be reduced to or incorporated into the vocabulary of physics, or of any of the other 'hard' sciences. one reason for this irreducibility springs from the fact that any correct account of an agent's beliefs and other attitudes must recognize the normative element inherent in cognitive contents. one attitude has logical relations with others; these relations, though distorted and deranged by the limits of our powers, nevertheless serve to locate, and thus identify, the contents of our thoughts. when we treat the world as mindless, as in the natural sciences, nothing corresponds to this dimension of the mental. a less often noticed property that sets our talk of mental states apart from the vocabulary of the advanced sciences is its dependence on causal concepts. ordinary physical talk, like psychological talk, is full of causal concepts: the notion of a catalyst is as causal as the notion of an intentional action. the difference lies in the promise, intrinsic to physics, but irrelevant to psychology, that the causal concept can, with time and research, be supplanted by an account of the mechanism which will explain what the clumsy causal notion merely finessed. the laws of physics may, if we please, be called causal. the point is that they do not employ causal concepts. the concentration of psychology on the causal role of reasons rules out any hope that the basic mental concepts can be fitted into a closed system of laws. there is one more (much debated) consideration which militates against the nomological or definitional reduction of mental concepts to those of physics, namely, the fact that propositional attitudes and related events and states are in part identified in terms of their causal and other relations to events extraneous in time and place to the agent they characterize. thus it is held, correctly in my opinion, that the history of an individual's learning and use of words and concepts is in central cases necessarily a factor in determining what the words mean, and the contents of the concepts. i also think interpersonal communication plays a necessary role in the possibility and nature of thought. if externalisms of these sorts are indeed dominant and unavoidable

' word and object, 219.

features of the mental, the impossibility of incorporating psychology into a unified scientific theory of the world is clear.

these remarks are no more than reminders of why many philosophers, including me, are convinced that psychology, as long as it is taken to include the concepts of intentional action, belief, perception, and the affective attitudes, cannot be made part of physics or any other 'natural' science. the question here is whether the definitional and nomological irreducibility of the prepositional attitudes is a reason to doubt the objective status of the propositional attitudes.

quine has said that he agrees with my token-token version of monism, which holds that each object or event identified in mental terms is identical with an object or event identifiable in physical terms, though the *classes* staked out by the mental vocabulary cannot be equated by definition or strict laws with classes definable in the physical vocabulary. this position, which 1 call anomalous monism, comes close to reconciling most, though not all, of quine's declarations about propositional attitudes. anomalous monism makes sense of the claim that attitudes are dispositions to behave in certain ways, which are in turn physiological states, which finally are physical states, as well as the claim that intentional descriptions are not reducible to behavioral or physical descriptions, and so are not suited to incorporation into any strict system of laws, the mental vocabulary is practical and indispensable, but it is not made for the most precise science. anomalous monism does not suggest that mental events and states are merely projected by the attributer onto an agent; on the contrary, it holds that mental events are as real as physical events, being identical with them, and attributions of states are as objective. quine's description of attitude attributions as dramatic portrayals does not imply that there is nothing to portray.

the fact that the mental vocabulary is not fit for inclusion in sciences like physics or physiology cannot in itself be taken as impugning the reality of the states, events, and objects it is used to describe. a perfected physics must comprehend every object and event, but this is an ontological and nomological requirement that defines the aim of physics: physics does not speak to the interests that demand other ways of characterizing things.

i have not yet touched directly on the issue of the indeterminacy of translation or interpretation. if the thesis of indeterminacy holds, doesn't this imply that propositional attitudes are less than real? quine may be thought to hint as much when he says 'brentano's thesis of the irreducibility of intentional idioms is of a piece with the thesis of indeterminacy of translation' since '[to] accept intentional usage at face value is ... to postulate translation relations as somehow objectively valid though indeterminate in principle relative to the totality of speech dispositions'.<sup>2</sup> doesn't this say that accepting the mental idiom means postulating something for which there is no empirical evidence — something that we therefore have no reason to believe exists?

in any case, whether or not quine's views entail that there is something unreal about the propositional attitudes, it is an idea common to a number of philosophers that indeterminism does undermine the reality of mental states. fodor and others have thought this constitutes a reductio of indeterminism. daniel dennett, on the other hand, endorses indeterminism, but agrees that it subtracts from the reality of mental states. i argue against the inference. i see no way around indeterminism, but think it leaves the reality of the mental untouched.

first, let me defend the passages from quine which i just quoted, and which seem to aver that 'intentional usage' has no empirical content. there is a confusion here that may in part be quine's fault, but it is easy to resolve. if by 'intentional usage' quine means talk of meanings and propositions as hypostatized by philosophers, it is true that he rejects these as without empirical or explanatory content. but it does not follow that ordinary attributions of attitudes, including interpretations of speech, are empty. indeed, how could they be if mental talk is, as quine allows, something we cannot do without?

there are, i think, two things that above all tempt us to suppose that there cannot be any indeterminacy about what we mean by our words or what we think. one is first person authority, the fact that there is a presumption that we know what we think in a way no one else can. the other is engendered by the semantics of the sentences that are used to attribute attitudes. let me take the second point first, and let me stick to belief sentences for simplicity. it is obvious that there is a potential infinity in the number of beliefs that might be assigned to someone, and it is clear that there is a potential infinity of sentences available, for we can create a well-formed sentence by putting any of an infinity of sentences after the words 'agnes believes that . . .'. but the only viable semantic devices we have

2 ibid. 221.

available to accommodate this situation involve taking 'believes' (or 'believes that') as a relational verb, and treating what follows as a singular term or description. the singular terms can't all be unstructured proper names; there aren't enough such. so if what follows 'believes' or 'believes that' is a singular term, a demonstrative device must be at work; otherwise we must have descriptions. but in either case, there must be an infinite store of *entities* to be picked out. once we have picked the appropriate entity, though, we have, as attitude attributers, assigned a content to the belief.

it is natural, perhaps, to suppose that the entity we refer to in order to specify the content of a belief is the *object* of the belief, an entity which the believer entertains or grasps, the *thing* the believer believes. given the relational semantics of sentences about beliefs (and, of course, the other attitudes), and the natural assumption that the entity that identifies the belief is the object grasped by the believer, we are virtually forced to conclude that the entity designated must be unique. for if more than one entity would do equally well in specifying the same state of belief, which entity should we say the believer grasps, or how can we say that the believer knows what he believes?

this conundrum can be solved by giving up the idea of treating the grammatical objects in belief sentences as terms that name psychologically real objects, objects known to or entertained by or grasped by the believer. the only *object* required for the existence of a belief is a believer. having a belief is not like having a favorite cat, it is being in a state; and being in a state does not require that there be an entity called a state that one is in. all that is necessary for the truth of an attitude attribution is that the predicate employed be true of the person with the attitude.

dummett deserves credit for characterizing antirealism in terms of truth rather than of reference, in terms of sentences rather than of names or descriptions. for the real issue about sentences that attribute mental states is not ontological; beliefs are not entities, nor do the 'objects of belief have to be objects. the real issue is whether or not attributions of attitudes are objectively true or false.<sup>3</sup>

the attribution of attitudes is analogous in many ways to the

<sup>&</sup>lt;sup>1</sup> there is, to be sure, the question whether meanings and propositions exist, hut this i see as a matter for decision. it they exist, they are abstract entities, and need no more than defining. once defined, it remains to be shown whether they do useful explanatory or descriptive work.

#### indeterminism and antirealism 75

measurement of various magnitudes. we can assign numbers to keep track of the sizes, weights, and speeds of objects provided the objects exhibit a pattern of the appropriate kind. we do not suppose there are empirical entities called weights or si/es or speeds that objects have. as carnap pointed out long ago, we should not think of the sentence 'this box weighs 8 pounds' as identifying two entities called 'the weight of this box' and '8 pounds', but rather as equating the weight of the box in pounds with the number 8. thus the ontology required consists of the objects that have weights, and the numbers. the numbers are not part of the weighty objects; they belong not to the empirical world but to us who need them in order to keep track of certain relations among objects. in the same way, the entities to which we relate thinkers when we attribute beliefs and other prepositional attitudes to them are not in the thinkers-not in their minds, or before their minds. having a belief is just exemplifying a property, having a certain predicate true of one; but in order to have enough predicates for all the beliefs we may wish to distinguish, we must construct the predicates by using a relational verb and filling in one of the places with a reference to an object drawn from some potentially infinite store. one such predicate is 'x believes that snow is white'.

one virtue of this analogy is that it makes clear why different assignments of objects can capture all the relevant information about a situation without compromising the truth or 'reality' of the situation. no one thinks the fact that we can register weight in either pounds or kilograms shows that there is something unreal about the weight of an object: different sets of numbers can be used to keep track of exactly the same facts. suppose the fact is that a weighs the same as b; then the number assigned to measure the weights of a and b is the same number, whether we measure weight in pounds or kilos. suppose the fact is that c weighs twice what d does; then the number assigned to measure the weight of c must be twice the number assigned to measure the weight of d, whether the weights are given in pounds or kilos. all that is necessary is that the entities we use to keep track of how much things weigh have a structure in which certain features of the objects weighed can be represented, and numbers can do this in endless ways.

another virtue of the analogy is that it makes clear why quine is justified when he insists, against chomsky and others, that the indeterminacy of translation is distinct from the underdetermination of

theories by all possible evidence. there is a question, to be sure, whether underdetermination in quine's sense can occur, for it requires that there be empirically equivalent, but incompatible, theories.<sup>4</sup> indeterminacy is not like this: the empirically equivalent theories it accepts as equally good for understanding an agent are not incompatible, any more than the measurement of weight in pounds or kilos involves incompatible theories of weight.

what objects can we use to keep track of the attitudes? clearly they must constitute approximately as complexly figured a field as the individual attitudes themselves. the most conspicuous features of the individual attitudes are their basically rational structures (if someone believes that everything is white, that person has a belief that entails that snow is white), and their relations to the world (the belief that snow is white is true if and only if snow is white).<sup>5</sup>

entities that have these required properties are our sentences, and it is not clear that any other set of entities will do as well (except utterances<sup>6</sup>). in any case, there is a sense in which our standard attributions of attitude do use the sentences of the attributer to specify the contents of the attributed attitude. by this i mean no more than that we must use a sentence of our own after such words as 'joan believes that . . .' to identify joan's belief. now we must ask which properties of the sentence are relevant to the identification of the belief, just as we ask which properties of numbers are essential to their role in recording weights. it is obvious that the property of having a truth value is not enough. the syntax would be more than enough if it were not for the fact that the same sentence can be used to mean quite different things, depending on the interpretation (the same sentence may belong to different languages or idiolects, or it may be ambiguous). because of this difficulty, i have suggested that we take the actual utterance (or inscription) provided by the attributer as the object to which he is referring to give the content of an attitude. so i proposed a paratactic account of attitude-attributing sentences, in which the 'that' of 'joan believes that . . .' is to be

<sup>&</sup>lt;sup>4</sup> quine has changed his mind on the issue more than once. see w. v. quine, *pursuit* of truth, 91 ff.

<sup>&</sup>lt;sup>5</sup> it should not be thought that by speaking of the relations between the attitudes and the world i am embracing a correspondence theory. a theory of truth of the sort i have in mind does depend on setting up a relation between certain words and objects (tarski's relation of 'satisfaction'), but it makes no use of objects to which sentences might corre spond.

<sup>&</sup>lt;sup>6</sup> see essay 4.

#### indeterminism and antirealism 77

taken as a demonstrative referring to the utterance that follows. the utterance has a unique syntax, and can be taken to resolve the question of the appropriate language and idiolect, to have eliminated ambiguities, and to have provided the parameters needed to fix indexical references. as a rough equivalent of the paratactic account, we could think of the content sentence as enclosed in quotation marks, with the language and other features of the speaker and context understood or made explicit. this second idea does better at capturing the intuition that attitude-attributing expressions are single sentences, but it is more complicated.

it should not surprise us that what we can say and understand about the propositional attitudes of others should be what we can capture by matching up our own sentences (or utterances; i shall not bother to distinguish) to those attitudes. this does not mean, of course, that someone else may not think things i cannot; but then that is bound to be something the content of which i cannot express either. 1 do not want, either, to rule out altogether propositional thoughts i have but cannot express. it is enough for my argument that when we can express a thought, our own or that of another, we must fall back on the basic device of representing it in the fabric of our sentences. our sentences provide the only measure of the mental.

what follows for indeterminacy? everyone is apt to agree that sometimes quite distinct sentences of the same reporter may be used to attribute exactly the same thought, so some degree of indeterminacy comes automatically with the redundancy of language. but the indeterminacies of which quine speaks, and most of which i accept, are another matter, for they suggest that sentences that no one would take to be even roughly synonymous may nevertheless be used to specify the same thought.

let me say briefly where quine and i differ and agree on these matters. both of us take as the basic evidence for interpretation a person's attitudes towards various utterances and the circumstances that cause these attitudes. quine takes certain patterns of assent and dissent to yield the interpretation of the truth-functional sentential functions, and so do 1; but i go on, as quine does not, to use the same method to locate the devices for quantification and cross-reference. thus i do not, like quine, see the internal structure of the simplest sentences as indeterminate. we both assume that the observed range of phenomena which prompt assent and dissent to certain sentences allow us to connect those sentences to events and objects in the

world. as for sentences less directly tied to the readily observable, i think we can often do better than quine believes, because i hold that we can use our ability to detect degrees of evidential support between sentences as a key to correct interpretation. there remain two important kinds of indeterminacy on which we agree: indeterminacy due to what quine calls the inscrutability of reference, and indeterminacy that results from the blurring of the distinction between the analytic and the synthetic.

the thesis of the inscrutability of reference says there is no way to tell which way of connecting words with things is the right way; if one way works, there will be countless others that do as well. from a technical point of view, this means that for the standard satisfaction relation (satisfaction is a sophisticated form of reference) we can substitute endless other relations without altering the truth conditions of any sentence or the logical relations among sentences. since all the evidence for interpreting language must come at the sentential level (for only sentences have a use in communication), the result is that there can be no evidence that one of the satisfaction (or reference) relations is the right one.

here is an example. suppose satisfaction relation s maps the word 'rome' onto rome, and the predicate 'is a city in italy' onto cities in italy. then the truth definition will show that the sentence 'rome is a city in italy' is true if and only if rome is a city in italy. now consider another satisfaction relation s' which maps the word 'rome' onto an area 100 miles to the south of rome, and the predicate 'is a city in italy' onto areas 100 miles south of cities in italy. the truth definition will now say that the sentence 'rome is a city in italy' is true if and only if the area 100 miles south of rome is an area 100 miles south of a city in italy. the truth conditions are clearly equivalent. the thesis of the inscrutability of reference contends that there can be no evidence that s is any better than s' for interpreting the sentence 'rome is a city in italy'. there is no telling what a sentence is 'about', or what someone is thinking about.

how can this be? surely rome and an area 100 miles to its south are not the same entity. true enough. however, any differences we conventionally think of as differences in the reference of names or the extensions of predicates will be preserved by any correct satisfaction relation. the fact that there is no empirical difference between the interpretations yielded by 5 and by s' doesn't entail that the person being interpreted can't tell the difference between rome

and an area 100 miles to the south; but it does entail that there is no saying which one his word 'rome' refers to. correct interpretation keeps track of a complex pattern, and locates particular sentences and attitudes within it. but the counters we use to represent this pattern, namely our sentences, can represent it in more than one way.

if the two interpretations of someone's utterance of 'rome is a city in italy' are equally correct, why is it that we feel that one is the right one? the complete answer is, i think, pretty complicated, but the simple answer is that we accept the standard, and in this case easiest, method of word-for-word translation. of course, when the homophonic translation manual is one of the available manuals, we will use it; and in interpreting a language for which no homophonic manual is available, we will use the conventional, and for the most part shortest, available method. this does not mean that other interpretations are wrong: s' provides as correct an interpretation of 'rome is a city in italy' as s.

to someone who objects, 'but "rome" doesn't *mean* "an area 100 miles south of rome" ', the right answer is: individual words don't have meanings. they have a *role* in determining the truth conditions of sentences, and this role is captured by *s* as surely as it is by .5.

still, this answer leaves something out, which is the necessity, in assigning contents to an agent's speech or attitudes, of sticking to some one method of interpretation. it will falsify 'rome is a city in italy' to interpret it as true if and only if rome is in an area 100 miles south of a city in italy; and it will destroy the entailment relations among sentences to use one method of interpretation for one sentence and another method for another sentence. just as we must indicate whether the numbers we are using to measure temperature place the temperature on the fahrenheit or the centigrade scale, so we must indicate which method of interpretation we are using (s or s', for example). one of the conveniences of sticking to the same method is that we can leave the indication of the method tacit. (somewhat similarly, if i say 'it must be 70 in this room', you will naturally understand me to be using the fahrenheit scale.)

there is a further objection which many people feel is fatal to the thesis of indeterminacy. the claim is that indeterminacy undermines first person authority. such people want to say, 'maybe quine's empirical criteria can't rule out the use of alternative methods of interpretation for others, but / know that i mean rome by my word "rome" and not the area 100 miles south of rome'. this remark is

correct, but it doesn't conflict with indeterminism. the reason it is correct is that the speaker does not identify rome and an area 100 miles to the south of rome. when a person is making such self-attributions, we know that words in the metalanguage play the same roles as those words do in the object language, for object language and metalanguage are one and the same. first person interpretations are necessarily tied to the homophonic translation manual (which is to say, translation, or interpretation, has no place here).

it should not be concluded from the fact that a person is restricted to a unique way of interpreting himself (if this can be called interpretation: it would be better to say that aside from pathological cases, our way of interpreting others has no application to ourselves) that therefore his words have unique reference. all that the agent has discovered in our example is that his phrases 'rome' and 'the area 100 miles south of rome' cannot be substituted for one another in his language, and this does not fix the reference of either expression. by the same token, the fact that the reference of his words is not fixed doesn't show that he doesn't know what he means or thinks. the different ways of interpreting his speech, or of representing his thoughts, mark no difference in the contents of his attitudes, so there is nothing he fails to know if he doesn't know which way of representing his thoughts is the right way; since there is no one right way; there is nothing more for him to know.

not all cases of indeterminacy depend on the possibility of systematically altering the satisfaction relation in ways that do not affect the truth conditions of sentences; there can be indeterminacies that affect truth values. these indeterminacies arise if one accepts quine's claim (as i do) that there is no principled way to make a clear distinction between the analytic and the synthetic. here it is easy to give examples. i find that i very often disagree with other people over whether to call the color of some object green or blue. the disagreement is consistent: there is a fairly definite range of cases where i say green and they say blue, we can account for this difference in two ways: it may be that i (or most other people) are wrong about the color of certain objects, or it may be that i don't use the words 'blue' and 'green' in quite the way others do. there may be no way to decide between these two accounts; by making compensatory adjustments elsewhere in one's interpretation of my sentences and beliefs one can accommodate either story, but on one account certain of my pronouncements about colors are false, while on the other they

are true. and on both accounts, i know what i think: i think the things i call green are green and the things i call blue are blue; neither account of what i mean and think threatens this first person knowledge.

where the criticisms i have been discussing suggest that if one accepts indeterminism one is treating meanings and mental states as less than totally real, daniel dennett has refreshingly complained that on my view intentional states are *too* real.<sup>7</sup> dennett sees clearly why my view, which compares rival theories of interpretation to different ways of representing length or temperature in the numbers, can't in itself call into question the reality of what is represented. but he thinks two different systems of belief attribution to an individual may differ substantially, even to the point of yielding different predictions of behavior, and yet nothing would establish that one system and not the other described the person's real beliefs.

here i think the issue of prediction is something of a red herring. no system of attitude attribution, no matter how complete, yields any prediction of actions without a theory, and it is certainly possible to differ on predictive theories in psychology. it is clear that no plausible predictions will be forthcoming without a quantitative description that specifies degrees of belief (subjective probabilities) and the relative strength of desires. but even supposing we could give a complete description of all the attitudes and their strengths for a given individual at a time, there is no reason to suppose that there exist strict laws to predict what the individual will do next—and good reasons to suppose there are no such laws. what sort of evidence can dennett appeal to, then, to show that equally justified systems for attributing attitudes are incompatible? the fact that different systems result in different predictions proves nothing, since the same system can be used to support different predictions.

dennett's idea is that what is real is behavior, and intentional states are patterns in this behavior. the patterns aren't defined in terms of the behavior—they are perceived by an observer when the observer takes the 'intensional stance'. the value of the patterns is that they reduce a vast, by us indescribably complex, physical situation to something we can grasp, and something on the basis of which we can make rough predictions. the patterns are in some (reduced) sense real, though abstract; and different people may

7 daniel dennett.'real patterns'.

perceive different patterns in the same behavioral field. some of these patterns may do better at predicting or understanding some phenomena; other patterns do better at predicting or understanding other phenomena. all these patterns are 'real', but when they are different, there is no saying which one represents the real attitudes of the targeted agent.

it is not easy to see how to judge this suggestion, and harder to grasp how it relates to the issue of realism. surely not every pattern one perceives in behavior is a pattern of propositional attitudes. but is any? patterns, dennett tells us, are abstractions, like centers of gravity, but then can they be beliefs and desires? beliefs and desires, we like to think, are states of a physical body which can have causal consequences; abstractions, 1 assume, have no causal relations. forces act not on centers of gravity, but on the things that have centers of gravity; but forces certainly alter my beliefs and intentions. do we perceive patterns? it seems to me not: what we perceive is something that has a certain pattern, and with luck (and the right stance) we may perceive *that* it has that pattern. so the question isn't whether patterns are real. being no nominalist, i think patterns, like shapes and numbers, are as real as can be. but i do not see how the propositional attitudes of a person can be patterns. if we ask what exhibits the pattern, we can say it is the person, or we can say it is the observable behavior of the person, but in either case no issue concerning the ontological status of attitudes is at stake. if people or their behavior really do exhibit the patterns dennett says they do, and for someone to have a propositional attitude is for that person to exhibit that pattern, then there is nothing ontologically or epistemically second-grade about the attitudes.

it seems to me dennett has confused two issues. one is whether the attitudes are entities, and here i think the answer is no, unless you suppose that states are entities. otherwise one should simply talk of people having attitudes, which means that certain predicates are true of them. the second issue is whether there is a correct answer to the question whether or not someone has a certain attitude. this i take to be, not a question about vagueness or borderline cases, but a question whether there are objective grounds for choosing among conflicting hypotheses. dennett has urged that the answer to the second question is that there are no such grounds; but i do not think he has given any reason to accept this answer.

are there objective grounds for choosing among conflicting

#### indeterminism and antirealism 83

hypotheses? especially in this case we have to ask what makes grounds 'objective'. the ultimate source (not ground) of objectivity is, in my opinion, intersubjectivity. if we were not in communication with others, there would be nothing on which to base the *idea* of being wrong, or, therefore, of being right, either in what we say or in what we think, the possibility of thought as well as of communication depends, in my view, on the fact that two or more creatures are responding, more or less simultaneously, to input from a shared world, and from each other. we are apt to say that someone responds in 'the same way' to, say, wolves. but of course, 'same' here means 'similar'. our grounds for claiming that a person finds one wolf similar to another is the fact that the person responds in similar ways to wolves. this prompts the next question: what makes the reactions similar? the only answer is, someone else finds both wolves and the reactions of the first person similar. this of course only puts the basic question off once more. nevertheless, it is this triangular nexus of causal relations involving the reactions of two (or more) creatures to each other and to shared stimuli in the world that supplies the conditions necessary for the concept of truth to have application. without a second person there is, as wittgenstein powerfully suggests, no basis for a judgement that a reaction is wrong or, therefore, right.8

this brings me to my last point. the analogy i proposed between measurement in the physical sciences and the assignment of contents to the words and thoughts of others is imperfect in an essential respect. in the case of ordinary measurement, we use the numbers to keep track of the facts that interest us. in the case of the propositional attitudes we use our sentences or utterances. but there is this difference: we can mutually specify the properties of the numbers. the numbers, like the objects we apply them to, lie, as it were, halfway between ourselves and others. this is what it means to say that they are objective or that they are objects. it cannot be this way with our sentences. you and i cannot come to agree on the interpretation of our sentences as a preliminary to using them to interpret others, for the process of coming to such an agreement involves interpretation of the very sort we thought to prepare for. it makes no sense to ask for a common standard of interpretation, for mutual interpretation provides the only standard we have.

<sup>&</sup>lt;sup>8</sup> for more on the subject of this and the concluding paragraphs, see essay 14.

we should not despair because we cannot provide a standard by which to judge the standard, a test for whether the standard meter bar is actually a meter long. our conclusion should rather be: if our judgements of the propositional attitudes of others are not objective, no judgements are, and the concept of objectivity has no application.

### 6 *the irreducibility of the concept of the self*

in his gentle, probing way dieter henrich often asked me whether i thought the idea of the self was an irreducible or primitive concept. i replied that i did think so, but so far as i remember, i never said why. i remain in doubt what an adequate response calls for, but here i shall try to list some of the considerations that may be relevant to a fuller answer.

a concept is irreducible only relative to specified resources. when it comes to the large, grand concepts that concern philosophers, like the good, truth, belief, knowledge, physical object, cause, and event, i think of a concept as irreducible if it cannot be defined in terms that are as general as the concept to be reduced, at least as clear, and that do not lead in a circle. with respect to the concepts i have listed, i think the search for such a definition or analysis is doomed. the question is, is the concept of the self another of the essential, and therefore irreplaceable, conceptual building blocks of our thought and languages?

i had better confess at once that i see no way to address this question directly, since the phrase 'the self doesn't play any clear role in ordinary speech. philosophers introduce it when they want to discuss such topics as self-consciousness, or the question what unites the various experiences of a person. so my approach will of necessity be oblique.

thomas nagel asks us (in 'the objective self) to imagine a description of the universe that tells when and where everything happens relative to some objectively given space-time framework; the description names each person with his or her mental states along with the positions and properties of all other objects. the description includes a complete history of nagel himself, all from a third person

point of view. the one thing it doesn't contain, nagel says, is the information that *he* is a particular one of the people described. of course, since he may not know his name is nagel, the only way he can with certainty locate himself in that world is by using the word t. thus the personal pronoun allows him to express knowledge that cannot be expressed in any other way. clearly, there is something irreducible, irreplaceable, embodied in the use of the first person pronouns—all the first person pronouns. in fact, all sentences containing indexicals depend, for their interpretation and their truth value, on who utters them. but if i utter them, i know without observation that it is i who uttered them. in this way, i relate myself to places, objects, times, and other people, by my use of 'there' ('here', 'behind me'), 'that' ('this'), 'now' ('tomorrow', all tensed verbs), 'you'. there is no substitute for this way of placing myself in the public world.

a less obvious, but equally important, function of indexicals is to provide an early and essential step into the domain of thought and language. for what we learn first is to associate what in the end turn out to be one-word sentences ('mama', 'no', 'dog', 'blue') with situations, events, objects, and their features. soon the child learns the magical power of making sounds adults find appropriate and hence reward. these are only preliminaries to fully fledged talk and thought, for in the beginning the child lacks awareness of the distinction between what is believed and what is the case, what is asked or demanded and what is answered or done. but though only preliminaries, these primitive relations between two people in the presence of stimuli from a shared world contain the kernel of ostensive learning, and it is only in the context of such interactions that we come to grasp the propositional contents of beliefs, desires, intentions, and speech.

i am not, of course, trying to establish the truth of these claims here; i have argued for them elsewhere. but if i am right, this view of the origin and nature of rational thought point to another sense in which the outlook on the world of each person differs irreducibly from that of any other person. the full force and meaning of this statement will emerge presently. but this much should be clear: the basic triangle of two people and a common world is one of which we must be aware if we have any thoughts at all. if i can think, i know that there are others with minds like my own, and that we inhabit a public time and space filled with objects and events many of which are (through the ostensions which made such thoughts available to us) known to others. in particular i, like every other rational creature, have three kinds of knowledge: knowledge of the objective world (without numerous successful ostensions, i would have no thoughts); knowledge of the minds of others; and knowledge of the contents of my own mind. none of these three sorts of knowledge is reducible to either of the other two, or to any other two in combination. it does indeed *follow* from the fact that i have any one of these sorts of knowledge that i have the other two since the basic triangle is a condition of thought, but none is conceptually or temporally prior to the others.

if my knowledge were exclusively of the contents of minds, perhaps of my own and that of others, there is no way i could construct my knowledge of the world we live in, for that requires the causal connections with the world provided by perception. of course, if i had knowledge of the contents either of my own mind or of another mind, i would have knowledge of the public world, and would know that i did. but this does not mean i could *construct* such knowledge from my knowledge of the contents of minds.

my knowledge of another mind is not reducible to my knowledge of the contents of my own mind and of the natural world, though my knowledge of the minds of others is dependent on my perceptual knowledge of the movements of their bodies. but to suppose that mental states can be *defined* in terms drawn exclusively from the natural sciences (as some behaviorists once held) is to suppose that the intentional is reducible to the extensional, and such reduction is surely not in the cards.

that knowledge of the contents of my own mind is special, and basic to all my knowledge, is, of course, part of the cartesian and empiricist dreams. and this much is correct: such knowledge is basic in the sense that without it i would know nothing (though selfknowledge is not sufficient for the rest), and special in that it is irreducibly different from other sorts of knowledge.

what unites the three varieties of knowledge, and makes their interdependence the grounding of all thought, is best illustrated by describing the process of ostension. consider the case of two people both of whom have developed languages. one asks, pointing, 'what's that?' 'a cormorant', the other answers. more questions and pointings may be needed before the learner gets it right, but astonishingly often shared habits of generalization do the work at one try.

here we have the three mature forms of knowledge at work. the learner and teacher each know what is in the other's mind, and a vital element in this understanding depends on the shared perceptual stimulus. more important, however, is the way this triangular set of relations reveals the basis of the objectivity of thought, its ability to latch on to something in the public world. in our imagined scenario, both participants already have the concept of objective content and of truth; they know what a physical object is, and what it is for a belief to be true or false. nevertheless, at the start of the ostension, the learner has nothing but a wordless, conceptless stimulus to connect with a concept and a word. the learner has, we are supposing, plenty of concepts available; but which is the right one? only the teacher can determine this. it takes two to triangulate the location of the distal stimulus, two to provide an objective test of correctness and failure. if this is so, it must be so from the start, before the learner has the idea of objectivity and is learning, by way of ostensions which are, to begin with, no more than ways of establishing associations, what it is to judge that a response is right or wrong, the possibility of thought comes with company. thus the gearing of verbal (and other) responses to situations, events, and objects through the prompting and perception of others plays a key role both in the acquisition of a first language, or the learning of a second language in the absence of an interpreter or bilingual dictionary.

learning a first and learning a second language are, of course, very different enterprises. the former is a matter of entering the domain of thought for the first time, the latter is a matter of someone already at home in the realm of thought entering into the thought of someone else. both, however, depend on similar mechanisms and similar cues. furthermore, the contrast is weakened by the realization that in the case of the child initiate, the two forms of learning mesh, for in absorbing the idea of an objective world, the child is simultaneously learning to communicate with others, which requires insight into the thoughts and intentions of those others.

making sense of the verbal and nonverbal behavior of others is not an enterprise limited to the special cases of learning first or further languages or adding new words to one's own. or we might say, treating the notion of a language more strictly than it is treated in ordinary talk, that all these are matters of learning a new language, since any addition to or alteration in our verbal resources makes ours a different, a new, language. if we look at the fine detail, no two idiolects are identical, nor is any person's language apt to remain unaltered for long. we are as interpreters constantly working out (often automatically) what others mean by what they say. we are also, both because of the interdependence of thought and speech and for obvious further reasons, constantly deciding what others think, intend, and want. much of the time we assume we have things right enough for the purposes at hand, and most of the time what small adjustments are needed come so easily that we are aware of no mental effort. it is only now and then that we realize we are coping with old words used in ways new to us, or that some piece of grammar cannot be treated as familiar. in such cases we become conscious that our interpretive skills are being tested. it is here that the unique contribution of the first person point of view—the 'self—becomes most apparent.

the conceptually developed thinker has two basic interpretive resources at his or her disposal in coming to understand the utterances and actions of others: the assumption of sufficient rationality to make these actions intelligible, and knowledge of how perception yields the contents of belief. in the case of speech, this is easy to illustrate, and the lesson carries over to the propositional attitudes. sentences, or rather the attitudes they express, owe their content, that is their meaning, to two things: their relations to other sentences or attitudes, and their relations, direct or indirect, to the world through perception. it is therefore impossible for an interpreter wholly to disregard the logical relations among a speaker's sentences or attitudes. this is not a matter of an effort on the part of agents to be consistent; it is a matter of their speech and behavior having the meaning they do because of how they hang together. unless there is sufficient coherence, it is impossible to assign propositional contents to their speech, beliefs, desires, or intentions. the interpreter's assumption of a degree of rationality on the part of those she wishes to understand is thus no more than a condition of understanding them at all.

the fundamental role of ostension in learning and interpreting speech guarantees that an interpreter cannot go wrong in generally supposing that a speaker's utterances reliably touched off by evident features of the observable world are true and about those features. mistakes on the part of speaker or interpreter are to be expected, but these cannot be the rule, since errors take their content from a background of veridical thought and honest assertion. the crucial difference

between the predominant, mostly banal, run-of-the-mill but on-target beliefs and assumptions and the occasional deviation is this: errors, confusions, irrationalities have particular explanations; getting things right, aside from hard cases, is to be expected.

1 have said little about knowledge of the contents of our own minds. like all knowledge, it cannot exist in isolation from its social beginnings; the concept of oneself as an independent entity depends on the realization of the existence of others, a realization that comes into its own with communication. but the vocabulary of the attitudes applies equally to oneself and to others, and the contents of the attitudes so attributed is likewise expressed in concepts that are in the public domain. my thought that shakespeare was a woman is mine, but anyone can think as i do. what, then, is so special about my knowledge of my own mind?

i not only know what i think. i also know an infinity of things that i can express and which i know someone else might believe or doubt or wonder about; the list is in a sense as large as the list of things expressible in my language, in concepts i command. these are the prepositional contents to which i advert when i attribute attitudes to myself or to others. this rich reservoir of conceptual resources is what i must use in interpreting the utterances or actions of those around me. when i speak of 'using' these resources, i do not mean it generally takes reflection or effort. it is done without notice, automatically. even your slips of tongue, omissions, fragmented grammar, accidental substitutions of one name for another, i silently, and often unconsciously, correct. most of the time, your words strike my ear and i understand them without any conscious intermediate mental process, but my understanding proves that a process, however inarticulate or far from introspective investigation, has taken place, a process that deserves to be called, if not interpretation, then something much like it. it is interpretation in which conscious reasoning and explicit recourse to evidence and induction have been reduced to zero.

interpretation in this etiolated sense can be described in terms, not of the observed process, but of the transition from input to outcome, for this transition is after all the same as the transition accomplished by conscious interpretation. here the elements are as before: the search for coherence based on the assumption of rationality, and the perception of external cues which are within the ken of the speaker or actor. my interest here is not in the details of the process, but in a

#### irreducibility of the concept of the self 91

tact so obvious it may escape notice: that the standards of rationality and reality on which i depend in understanding others are my own, and there can be no appeal beyond them. this is not to deny that what standards [ have would not exist if it were not for a history of communication and experience. nor is it to suppose i cannot reflect on my own reasoning and consult with others for greater clarity, wisdom, and information. but insofar as i seek information directly by experiment and observation, i again can do no better or more than employ my own resources. and if i wonder whether the norms of rationality i employ in trying to comprehend others are correct, i can, of course, ask sebastian whether i am as objective or reasonable as i should be in my account of basil's thoughts and actions. but my understanding of sebastian's reply will be one more exercise of my own standards and methods. there is another obvious indication of the irreducible singularity of my direct acquaintance with the contents of my own mind, and this is that such knowledge is unique in that it is, aside from unusual cases, unsupported by observation, evidence, or reasons. this is due, at least in part, to the fact that here interpretation has no application. drawing on my store of potential thoughts will yield nothing but tautologies when applied to those same thoughts. self-consciousness of this kind can direct attention inward and promote self-criticism, but it can lead from insight to action only indirectly.

i do not think that, because the ultimate court of appeal is personal, therefore my judgements are arbitrary or subjective, for they were formed in a social nexus that assures the objectivity if not the correctness of my beliefs. intersubjectivity is the root of objectivity, not because what people agree on is necessarily true, but because intersubjectivity depends on interaction with the world. though we could not have been at the point of comparing notes without prior interaction, it is private notes that in the end get compared. it is here that each person, each mind or self, reveals itself as part of a community of free selves. there would be no thought if individuals did not play the indispensable, and ultimately unavoidably creative, role of final arbiter. intersubjective

# 7 rational animals

neither an infant 1 week old nor a snail is a rational creature. if the infant survives long enough, he or she will probably become rational, while this is not true of the snail. if we like, we may say of the infant from the start that he is a rational creature because he will probably become rational if he survives, or because he belongs to a species with this capacity. whichever way we talk, there remains the difference, with respect to rationality, between the infant and the snail on one hand, and the normal adult person on the other.

the difference consists in having propositional attitudes such as belief, desire, intention, and shame. this raises the question how to tell when a creature has propositional attitudes; snails, we may agree, do not, but how about dogs or chimpanzees? the question is not entirely empirical, for there is the philosophical question what evidence is relevant to deciding when a creature has propositional attitudes.

some animals think and reason; they consider, test, reject, and accept hypotheses; they act on reasons, sometimes after deliberating. imagining consequences. and weighing probabilities; they have desires, hopes, and hates, sometimes for good reasons. they also make errors in calculation, act against their own best judgement, or accept doctrines on inadequate evidence. any one of these accomplishments, activities, actions, or errors is enough to show that such an animal is a rational animal, for to be a rational animal is just to have propositional attitudes, no matter how confused,

contradictory, absurd, unjustified, or erroneous those attitudes may be. this, i propose, is the answer. the question is: what animals are

the question is: what animals are rational? of course i do not intend to name names, even names of species or other groups. i shall
not try to decide whether dolphins, apes, human embryos, or politicians are rational, or even whether all that prevents computers from being rational is their genesis. my question is what makes an animal (or anything else, if one wants) rational.

the propositional attitudes provide an interesting criterion of rationality because they come only as a matched set. obviously a rich pattern of beliefs, desires, and intentions suffices for rationality; yet it may seem far too stringent to make this a necessary condition. but in fact the stringency lies in the nature of the propositional attitudes, since to have one is to have a large complement. one belief requires many beliefs, and beliefs demand other basic attitudes such as intentions, desires, and, if i am right, the gift of tongues. this does not mean that there are not borderline cases. nevertheless, the intrinsically holistic character of the propositional attitudes makes the distinction between having any and having none dramatic.

to make the distinction so strong, and to make it depend on language, invites an accusation of anthropocentrism. the complaint is just, but it ought not to be leveled against me. i merely describe a feature of certain concepts. after all, it is not surprising that our human language is rich in resources for distinguishing men and women from other creatures, just as the inuit are said to have a vocabulary convenient for picking out varieties of snow (this is now said to be a myth). we connive with our language to make it, and us, seem special.

i promised not to discuss the question whether particular species are rational, but it will be impossible to avoid the appearance of talking of the feats and abilities of beasts because so much discussion of the nature of thought has by tradition centered on the mental powers of nonhuman animals. i consider this approach as just a colorful (and sometimes emotionally laden) way of thinking about the nature of thought.<sup>1</sup>

norman malcolm tells this story, which is intended to show that dogs think:

<sup>&</sup>lt;sup>1</sup> my records tell me that before this essay was written i gave no less than ten talks, from valdosta, georgia, to auckland, with the title 'why animals can't think'. the title was tendentious, since what 1 argued for (as here) was that only creatures with a language can think. i happen to believe, however, that men and women are alone in having language, or anything enough like a language to justify attributing propositional thoughts to them. on the moral issue how we should treat dumb creatures, i see no reason to be less kind to those without thoughts or language than to those with; on the contrary.

suppose our dog is chasing the neighbor's cat. the latter runs full tilt toward the oak tree, but suddenly swerves at the last moment and disappears up a nearby maple. the dog doesn't see this maneuver and on arriving at the oak tree he rears up on his hind feet, paws at the trunk as if trying to scale it, and barks excitedly into the branches above. we who observe this whole episode from a window say, 'he thinks that the cat went up that oak tree'.<sup>2</sup>

(malcolm added, we would say the dog was barking up the wrong tree.) malcolm claims that under the circumstances someone who attributed that belief to the dog might well—almost surely would— be right; he would have exactly the sort of evidence needed to justify such an attribution.

let me give a preliminary argument designed to put malcolm's claim in doubt. it's clear that the evidence for the dog's 'belief depends on taking belief as a determinant of action and emotional response. we are asked to infer from what we see that the dog wants to catch the cat, that he runs where he does because of this desire and a belief about where the cat has gone, and that he is venting his frustration at not being able to follow the cat up the tree by barking, pawing the ground, and so forth. the details do not need to be right, of course. the point is so far obvious: if we are justified in inferring beliefs, we are also justified in inferring intentions and desires (and perhaps much more).

but how about the dog's supposed belief that the cat went up that oak tree? that oak tree, as it happens, is the oldest tree in sight. does the dog think that the cat went up the oldest tree in sight? or that the cat went up the same tree it went up the last time the dog chased it? it is hard to make sense of the questions. but then it does not seem possible to distinguish between quite different things the dog might be said to believe.

one way of telling that we are attributing a propositional attitude is by noting that the sentences we use to do the attributing may change from true to false if, in the words that pick out the object of the attitude, we substitute for some referring expression another expression that refers to the same thing. the belief that the cat went up that oak tree is not the same belief as the belief that the cat went up the oldest tree in sight. if we use words like 'believe', 'think', and 'intend', while dropping the feature of semantic opacity, there is a question whether we are using those words to attribute propositional

<sup>2</sup> norman malcolm.'thoughtless brutes', 13.

attitudes. for it has long been recognized that semantic opacity distinguishes talk about propositional attitudes from talk of other things.

someone may suggest that the position occupied by the expression 'that oak tree' in the sentence 'the dog thinks the cat went up that oak tree' is, in quine's terminology, transparent, the right way to put the dog's belief (the suggestion continues) is 'the dog thinks, with respect to that oak tree, that the cat went up it' or that oak tree is the one the dog thinks the cat went up'. but such constructions, while they may relieve the attributer of the need to produce a description of the object that the believer would accept, nevertheless imply that there is some such description; the *de re* description picks out an object the believer could somehow pick out. in a popular if misleading idiom, the dog must believe, under some description of the tree, that the cat went up that tree. but what kind of description would suit the dog? for example, can the dog believe of an object that it is a tree? this would seem impossible unless we suppose the dog has many general beliefs about trees: that they are growing things, that they need soil and water, that they have leaves or needles, that they burn. there is no fixed list of things someone with the concept of a tree must believe, but without many general beliefs, there would be no reason to identify a belief as a belief about a tree, much less an oak tree. similar considerations apply to the dog's supposed thinking about the cat.

we identify thoughts, distinguish among them, describe them for what they are, only as they can be located within a dense network of related beliefs. if we really can intelligibly ascribe single beliefs to a dog, we must be able to imagine how we would decide whether the dog has many other beliefs of the kind necessary for making sense of the first. it seems to me that no matter where we start, we very soon come to beliefs such that we have no idea at all how to tell whether a dog has them, and yet such that, without them, our confident first attribution looks shaky.

not only does each belief require a world of further beliefs to give it content and identity, but every other propositional attitude depends for its particularity on a similar world of beliefs. in order to believe the cat went up the oak tree i must have many true beliefs about cats and trees, this cat and this tree, the place, appearance, and habits of cats and trees, and so on; but the same holds if i wonder whether the cat went up the oak tree, fear that it did, hope that it did, wish that it had, or intend to make it do so. belief—indeed, true belief—plays a central role among the propositional attitudes. so let me speak of all the propositional attitudes as thoughts.

as remarked above, there may be no fixed list of beliefs on which any particular thought depends. nevertheless, much true belief is necessary. some beliefs of the sort required are general, but plausibly empirical, such as that cats can scratch or climb trees. others are particular, such as that the cat seen running a moment ago is still in the neighborhood. some are logical. thoughts, like propositions, have logical relations. since the identity of a thought cannot be divorced from its place in the logical network of other thoughts, it cannot be relocated in the network without becoming a different thought, radical incoherence in belief is therefore impossible, to have a single propositional attitude is to have a largely correct logic, in the sense of having a pattern of beliefs that logically cohere. this is one reason why to have propositional attitudes is to be a rational creature, the point extends to intentional action, intentional action is action that can be explained in terms of beliefs and desires whose propositional contents rationalize the action. similarly, an emotion like being pleased that one has stopped smoking must be an emotion that is rational in the light of beliefs and values one has.

this is not to deny the existence of irrational beliefs, actions, and emotions, needless to say. an action one has reasons to perform may be an action one has better reasons to avoid. a belief may be reasonable in the light of some but not the totality of one's other beliefs; and so on. the point is that the possibility of irrationality depends on a large degree of rationality. irrationality is not mere lack of reason but a disease or perturbation of reason.

i assume that an observer can under favorable circumstances tell what beliefs, desires, and intentions an agent has. indeed, i appealed to this assumption when i urged that if a creature cannot speak, it is unclear that intensionality can be maintained in the descriptions of its purported beliefs and other attitudes. similarly, i wondered whether, in the absence of speech, there could be adequate grounds for attributing the general beliefs needed for making sense of any thought. without defending the assumption that we can know other minds, let me distinguish this assumption from other stronger assumptions. merely to claim that an observer can under favorable conditions tell what someone else is thinking is not to embrace verificationism, even with respect to thoughts. for the observability

assumption does not imply that it is possible to state explicitly what evidence is necessary or sufficient to determine the presence of a particular thought; there is no suggestion that thinking can somehow be reduced definitionally to something else. nor does the observability assumption imply that the only way to determine the existence of a thought is by observing. on the contrary, it is clear that people normally know without observation or evidence what they believe, want, and intend.

nor does the observability assumption amount to behaviorism. propositional attitudes can be discovered by an observer who witnesses nothing but behavior without the attitudes being in any way reducible to behavior. there are conceptual ties between the attitudes and behavior which are sufficient, given enough information about actual and potential behavior, to allow correct inferences to the attitudes.

from what has been said about the dependence of beliefs on other beliefs, and of other propositional attitudes on beliefs, it is clear that a very complex pattern of behavior must be observed to justify the attribution of a single thought. or, more accurately, there has to be good reason to believe there is such a complex pattern of behavior. and unless there is actually such a complex pattern of behavior, there is no thought.

i think there is such a pattern only if the agent has language. if this is right, then malcolm was justified in attributing thought to his dog only if he believed, on good evidence, that his dog had language.

the view that thought—belief, desire, intention, and the like requires language is controversial, but certainly not new. the version of the thesis which i want to promote needs to be distinguished from various related versions. i don't, for example, believe that thinking can be reduced to linguistic activity. i find no plausibility in the idea that thoughts can be nomologically identified with, or correlated with, phenomena characterized in physical or neurological terms. nor do i see any reason to maintain that what we can't say we can't think. my thesis is not, then, that each thought depends for its existence on the existence of a sentence that expresses that thought. my thesis is rather that a creature cannot have a thought unless it has language. in order to be a thinking, rational creature, the creature must be able to express many thoughts, and above all, be able to interpret the speech and thoughts of others.

as i remarked above, this has often been claimed; but on what

grounds? given the popularity of the idea, from the rationalists through the american pragmatists, and even among contemporary analytic philosophers, there is a remarkable dearth of arguments. so far, i have pointed to the dubious applicability of the intensionality test where dumb animals are concerned, and the requirement, if thought is to be present, that there be a rich supply of general (and true) beliefs. these considerations point in the direction of language, but they do not amount to a demonstration that language is necessary to thought. indeed, what these considerations suggest is only that there probably can't be much thought without language.

against the dependence of thought on language is the plain observation that we succeed in explaining, and sometimes predicting, the behavior of languageless animals by attributing beliefs and desires and intentions to them. this method works for dogs and frogs much as it does for people. and, it may be added, we have no general and practical alternative framework for explaining animal behavior. don't these facts amount to *a justification* of the application of the method?<sup>3</sup>

no doubt they do. but there could remain a clear sense in which it would be wrong to conclude that dumb (= incapable of interpreting or engaging in linguistic communication) animals have propositional attitudes. to see this it is only necessary to reflect that someone might easily have no better or alternative way of explaining the movements of a heat-seeking missile than to suppose the missile wanted to destroy an airplane and believed it could by moving in the way it was observed to move. this uninformed observer might be justified in attributing a desire and beliefs to the missile; but he would be wrong. i know better, for example, not because i know how the missile is designed, but because i know that it moves as it does because it was designed and built by people who had the very desire and beliefs my ignorant friend assigned to the missile. my explanation, while still teleological, and dependent on the existence of propositional attitudes, is a better explanation because it does not attribute to the missile the potentiality for the rich range of behavior that a thinking creature must have.

the case of a languageless creature differs from the case of the missile in two respects: many animals are far more like humans in the range of their behavior than missiles are: and we often do not

<sup>&</sup>lt;sup>3</sup> this is the position stressed by jonathan bennett. *linguistic behavior*.

have a better way of explaining their behavior than by appropriating prepositional attitudes. what we need, then, in order to make a case, is a characterization of what it is that language supplies that is necessary for thought. for if there is such a necessary condition, we can continue to explain the behavior of speechless creatures by attributing propositional attitudes to them while at the same time recognizing that such creatures do not actually have propositional attitudes. we will be bound to acknowledge that we are applying a pattern of explanation that is far stronger than the observed behavior requires, and to which the observed behavior is not subtle enough to give point.

in the rest of this essay i state the condition for thought that i believe only language can supply, and i marshal considerations in favor of my view. although i present these considerations as an argument, it will be clear that my reasoning can be challenged at several points.

the 'argument' has two steps. i think i have shown that all the propositional attitudes require a background of beliefs, so i shall concentrate on conditions for belief. without belief there are no other propositional attitudes, and so no rationality as i have characterized it.

*first,* i argue that in order to have a belief, it is necessary to have the concept of belief.

*second*, i argue that in order to have the concept of belief one must have language.

norman malcolm, in the article mentioned above, makes a distinction similar to the one i want between having a belief and having the concept of a belief, but his terminology differs from mine. i have been using the word 'thought' to cover all the propositional attitudes. malclom, however, restricts the application of 'thought' to a higher level of thinking. in his view, the dog can believe the cat went up that oak tree, but it cannot have the thought that the cat has gone up that oak tree. the latter, but not the former, malcolm holds, requires language. malcolm makes the distinction by saying a creature merely thinks (believes) that p if it is aware that p, but it has the thought that p if it is aware that it is aware that p and believing that one believes that p. the second is a belief about a belief, and so requires the concept of belief. to make a rough comparison:

malcolm holds that language draws a line between creatures that merely think and creatures that have the concept of a thought; i hold that in order to think one must have the concept of a thought, and so language is required in both cases.

donald weiss takes issue with malcolm: weiss thinks it is possible to make sensible attributions of awareness to speechless creatures.<sup>4</sup> since i think his example may strike a responsive chord in others, let me paraphrase and then quote him at some length. here is the story. arthur is not a dog, but, let us say, a superdog from another planet. arthur arrives on earth unaccompanied, and here he hatches. he has no commerce with, or knowledge of, other creatures—he is observed through one-way mirrors. he has no language. according to weiss, we become convinced he has reflective intelligence when we witness this scene:

one day arthur comes upon a shiny metal, puts it in the fire, tries to hammer it out—but discovers that it is apparently no more malleable than it was when cold. he tries again more slowly and more methodically—but again the same result. the regularity in which arthur believed—we whisper among ourselves—is not entirely universal. arthur has discovered an instance that does not conform to the general rule.

arthur proceeds to walk agitatedly around his living space. he abruptly sits down; just as abruptly he gets up again; he paces forward and back. once more he sits down, but this time he remains seated. fifteen minutes pass without change of posture; arthur's eyes are focused straight ahead. then suddenly he leaps up and immediately proceeds to pile a large quantity of wood onto his fire ... he then plunges his newly discovered metal into the fire, and, after a time, withdraws it. he again attempts to hammer it out—and this time he meets with success. thus apparently satisfied ... he proceeds in a leisurely manner to cook himself a meal.<sup>5</sup>

weiss says we now have strong evidence arthur has reflected upon his own beliefs; he is particularly impressed by the fact that arthur in response to his state of befuddlement, sits wide-eyed and stock-still, and then veritably leaps to perform the acts that constitute the solution to his problem.<sup>6</sup>

i will ignore the question-begging vocabulary weiss uses in describing arthur's movements, for i think weiss is barking up a right tree: it is essential that we be able to describe arthur as being surprised. what i think is clear is that if he is surprised, he does have reflective thoughts, and, of course, beliefs.

<sup>&</sup>lt;sup>4</sup> donald weiss, "professor malcolm on animal intelligence'. <sup>5</sup> ibid., 91-2. <sup>6</sup> ibid.

this is not to claim that all thinking is self-conscious, or that whenever we think that p we must be aware that p, or believe that we believe that p, or think that we think that p. my claim is rather this: in order to have any propositional attitude at all, it is necessary to have the concept of a belief, to have a belief about some belief. but what is required in order to have the concept of a belief? here i turn for help to the phenomenon of surprise, since i think that surprise requires the concept of a belief.

suppose i believe there is a coin in my pocket. i empty my pocket and find no coin. i am surprised. clearly enough i could not be surprised (though i could be startled) if i did not have beliefs in the first place. and perhaps it is equally clear that having a belief, at least one of the sort 1 have taken for my example, entails the possibility of surprise. if i believe i have a coin in my pocket, something might happen that would change my mind. but surprise involves a further step. it is not enough that i first believe there is a coin in my pocket, and after emptying my pocket i no longer have this belief. surprise requires that i be aware of a contrast between what i did believe and what i come to believe. such awareness, however, is a belief about a belief: if i am surprised, then among other things i come to believe that my original belief was false. i do not need to insist that every case of surprise involves a belief that a prior belief was false (though 1 am inclined to think so). what i do want to claim is that one cannot have a general stock of beliefs of the sort necessary for having any beliefs at all without being subject to surprises that involve beliefs about the correctness of one's own beliefs. surprise about some things is a necessary and sufficient condition of thought in general. this concludes the first part of my 'argument'.

much of the point of the concept of belief is that it is the concept of a state of an organism which can be true or false, correct or incorrect. to have the concept of belief is therefore to have the concept of objective truth. if i believe there is a coin in my pocket, i may be right or wrong; i'm right only if there is a coin in my pocket. if i am surprised to find there is no coin in my pocket, i come to believe that my former belief did not correspond with the state of my finances. i have the idea of an objective reality which is independent of my belief.

a creature may interact with the world in complex ways without entertaining any propositions. it may discriminate among colors, tastes, sounds, and shapes. it may learn, that is change its behaviour, in ways that preserve its life or increase its food intake. it may 'generalize', in the sense of reacting to new stimuli as it has come to react to prior stimuli. yet none of this, no matter how successful by my standards, shows that the creature commands the contrast between what is believed and what is the case, as required by belief.

what *would* show command of this contrast? clearly linguistic communication suffices. to understand the speech of another, i must be able to think of the same things she does; i must share her world. i don't have to agree with her in all matters, but in order to disagree we must entertain the same propositions, with the same subject matter, and the same concept of truth. communication depends on each communicator having, and correctly thinking that the other has, the concept of a shared world, an intersubjective world. but the concept of an intersubjective world is the concept of an objective world, a world about which each communicator can have beliefs.

i suggest, then, that the concept of intersubjective truth suffices as a basis for belief and hence for thoughts generally. and perhaps it is plausible enough that having the concept of intersubjective truth depends on communication in the full linguistic sense. to complete the 'argument', however, i need to show that the *only* way one could come to have the belief-truth contrast is through having the concept of intersubjective truth. i confess i do not know how to show this. but neither do i have any idea how else one could arrive at the concept of an objective truth. in place of an argument for the first step, i offer the following analogy.

if i were bolted to the earth, i would have no way of determining the distance from me of many objects. i would only know they were on some line drawn from me towards them. i might interact successfully with objects, but i could have no way of giving content to the question where they were. not being bolted down, i am free to triangulate. our sense of objectivity is the consequence of another sort of triangulation, one that requires two creatures. each interacts with an object, but what gives each the concept of the way things are objectively is the base line formed between the creatures by language. the fact that they share a concept of truth alone makes sense of the claim that they have beliefs, that they are able to assign objects a place in the public world.

the conclusion of these considerations is that rationality is a social trait. only communicators have it.

# 8 the second person

**...meaning** something is like going up to someone.

(wittgenstein, philosophical investigations, §457)

how many competent speakers of a language must there be if anyone can be said to speak or understand a language? since this is a matter governed by the crooked course of evolution, i have no idea what the answer is; perhaps it takes quite a crowd. but as philosophers we can ask the question in a more theoretical vein. in this essay i shall be concentrating on the role-the role in principle-of the second person. my subject is not, i should perhaps add, the grammatical second person, the 'you' or 'thou', the 'tu' or 'vosotros'; i shall be writing about real second people, not the words used to address them.

a language may be viewed as a complex abstract object, defined by giving a finite list of expressions (words), rules for constructing meaningful concatenations of expressions (sentences), and a semantic interpretation of the meaningful expressions based on the semantic features of individual words, i shall not be concerned in this essay with the details of how such objects should be described or defined.

thought of this way, a language is abstract in the obvious sense that it is unobservable, changeless, and its components are also unob-servable and changeless. expressions may, if we wish, be viewed as acoustical or twodimensional geometric shapes that could, on occasion, inform actual utterances or inscriptions, but expressions themselves would remain abstract and their existence independent of exemplification.<sup>1</sup> the functions that

# interpret some expressions by

<sup>1</sup> or. to take equally abstract entities, expressions may he thought of as classes of utterances or inscriptions. but if we take expressions to be classes of utterances or inscriptions, all unuttered and unwritten expressions, and hence all expressions of unused languages, will be identical. in some cases there are awkward ways around this. (see w. v. quine and

mapping them on to objects or classes of objects are also, of course, abstract. the only concrete particulars that enter into the characterization of a language are some of the objects on to which some expressions are mapped (for example, by the naming relation, or, at another remove, by tarski's satisfaction relation).

the abstract character of language is nothing to wonder at, the concept of a language is of a sort with, and depends on, concepts like name, predicate, sentence, reference, meaning, with observation. these are all theoretical concepts, and the items to which they apply are abstract objects. we don't need them in order to use or learn a language; obviously they are not available to us when we are learning a first language. where we want these concepts is in talking about speech behavior. philosophers, psychologists, and linguists need these theoretical terms if they want to describe, theorize about, and explain verbal activities. the rest of us also have occasion to talk about talk, or write about writing, so these theoretical concepts have their place in the loose informal 'theories' we all have about language, indeed, we all talk so freely about language, or languages, that we tend to forget that there are no such things in the world; there are only people and their various written and acoustical products. this point, obvious in itself, is nevertheless easy to forget, and it has consequences that are not universally recognized.

a feature of the concept of a language as i have described it is, then, that there must be an infinity of 'languages' no one ever has spoken or ever will speak. to say someone speaks a particular language, say french, is just to say that his or her datable utterances and writings are tokens of french expressions. to be the token of a french sentence an utterance or inscription must instantiate a french sentence; that is, have one of the shapes defined to be a french sentence; and the utterance must have the semantic features the definition of french assigns to the shape. (other, probably unspoken, languages will assign other semantic features to these same shapes.) the existence of the french language does not depend on anyone's speaking it, any more than the existence of shapes depends on there being objects with those shapes.<sup>2</sup>

nelson goodman, 'steps toward a constructive nominalism'.) for our purposes it will be better to take expressions to be shapes, i.e. properties that utterances and inscriptions can have.

 $^{\rm 2}$  this concept of a language is essentially that of david lewis, 'languages and language'.

it follows that there is nothing about the existence of a particular language that imbues it with anything more than the sort of interest any abstract object may have; as logicians we can study it as one example among countless others of a formal pattern. such research pays dividends of various sorts, but it is unrelated to our normal concern with understanding the speech of others, or learning to make ourselves understood by them. our practical, as opposed to our purely theoretical, interest in linguistic phenomena is this: we want to understand the actual utterances of others, and we want our utterances to be understood. what has language to do with this interest?

the answer is that it is only by employing such concepts as word and sentence that we can give a systematic description of the linguistic aspects of linguistic behavior and aptitudes. we could not, for example, say what we have learned when we learn that 'demain' means tomorrow in french if we could not speak of words—those mysterious abstract (acoustical) shapes which utterances of the word 'demain' share. thus an utterance of the word 'demain' refers to the day after the day of the utterance. there is no easy way we could specify which utterances constitute utterances of sentences, and therefore constitute intelligible utterances, if we could not refer to words and sentences.

the main point of the concept of a language, then, and its attendant concepts like those of predicate, sentence, and reference, is to enable us to give a coherent description of the behavior of speakers, and of what speakers and their interpreters know that allows them to communicate. i do not suggest that speakers and those who understand them must themselves be able to provide such descriptions of their abilities and behavior, to illustrate: a competent speaker of a language (and a competent interpreter) knows the truth conditions of an indefinitely large number of sentences. so most english speakers know that an utterance of the sentence 'montreal is in canada' is true if and only if montreal is in canada; and the speaker knows an analogous fact about endless other sentences. the speaker doesn't need to put this knowledge into words. but we cannot describe the totality of this knowledge possessed by the speaker or interpreter of the language without ourselves having a theory-a theory of truth, or something like it—which is part of the description of english. (this description of what an english speaker knows does not have to be stated in english, and if it is not, it does not sound so trivial.)

to return now to the question with which i began: how many

speakers or interpreters of a language must there be for there to be one speaker? let me start with an apparent difficulty. to speak a language, one's utterances must be consistent with the definition of some language. the trouble is that utterances are finite in number, while the definition of a language assigns meanings to an infinite number of sentences. there will therefore be endless different languages which agree with all of a speaker's actual utterances, but differ with respect to unspoken sentences. what makes the speaker a speaker of one of these languages rather than another? and the problem may be worse still. for even if a speaker were (impossibly) to utter every sentence in some one language, many other languages would be consistent with all his behavior and internal states, as quine has maintained; and i agree.

the fact that all the publicly available evidence with regard to a speaker or group of speakers, even if imagined to exhaust all possible such evidence, might be consistent with many different languages (in the sense of 'language' we have temporarily ordained) ought not in itself to worry us, for we can agree that it is enough to know that a speaker speaks any one of a set of empirically equivalent languages, as long as the empirical constraints clearly define the set. so let us henceforth simply call any language in this set 'the language' of a speaker.3 this strategy is good enough for empirically equivalent languages if the evidence is imagined to contain an utterance of every sentence we would count as belonging to the language. but of course such evidence is never available. so there will be endless languages consistent with all the actual utterances of a speaker none of which is 'the' language the speaker is speaking.

the problem can be stated in a temporal mode, and addressed to an interpreter. if you (the interpreter) do not know how a speaker is going to go on, you do not know what language she speaks, no matter how much she has said up until now. it will not help to mention that the speaker has performed according to expectation so far, or that she went to the same school you did, or belongs to the same culture or community, for the question concerns not the past but the future. nor can we appeal to the idea that the speaker has mastered a set of

<sup>3</sup> i regard the existence of empirically equivalent languages (that is, languages equally consistent with all possible empirical evidence) as no more threatening to the reality or objectivity of the correct interpretation of utterances and their accompanying mental states than the existence of various scales for recording temperatures or lengths is to the reality or objectivity of temperature or length. (see essay 5.)

conventions (which conventions?), or has learned a set *of* rules (which ones?). the concepts of conventions or rules, like the concept of a language, cannot be called on to justify or explain linguistic behavior; at best these concepts help describe (i.e. define) linguistic behavior.

this difficulty, though it may have troubled wittgenstein, and certainly troubled kripke,<sup>4</sup> seems to me to have a relatively simple answer, the longer we interpret a speaker with apparent success as speaking a particular language, the greater our legitimate confidence that the speaker is speaking that language-that is, that she will continue to be interpretable as speaking that language. our strengthening expectations are as well founded as our evidence and ordinary induction make them. these expectations are in the main conditional. we do not usually know what someone will say, but we are prepared to interpret any of a very large number of things that person might say. our dispositions to interpret and a speaker's dispositions to go on in certain ways are not shadowy or mysterious: they are real features of brains and muscles. of course our beliefs about what is true of another person, and therefore what we expect that person to mean by what he does or might .say, may easily be wrong. i think such beliefs are frequently wrong, but far more often they are right, and the things we are right about usually put us in a position to correct our understanding of an utterance which does not belong to a language we thought was being spoken. to the extent that we are right about what is in someone's head, and therefore are right about what he would mean by endless things he does not say, we are right about 'the' language he speaks.

this very partial answer to the question what reasons an interpreter can have for believing that a speaker is speaking one language rather than another language that is equally compatible with the speech behavior that the interpreter has observed does not depend heavily on the details of how we explain successful interpretation. the point of the answer is that there are not *two* questions, one about reasons for believing a speaker is speaking one language rather than another, and a second question about how we naturally form expectations; the first question is simply a case of the second.

there is another aspect of interpretation, however, that is essential to our concerns: an interpreter (correctly) interprets an utterance of a

<sup>&</sup>lt;sup>4</sup> saul kripke. wittgenstein on rules and private language.

speaker only if he knows that the speaker intends the interpreter to assign certain truth conditions to his (the speaker's) utterance.<sup>5</sup> a full account of this thesis would require an explanation of the idea of 'assigning truth conditions' to an utterance, and this idea is no doubt as difficult to understand in relevant respects as the concept of meaning itself. but my aim here is not to solve that problem. it is only to emphasize, following grice, the central importance of intention in communication. if, with grice, we were sure that in order to mean something a speaker must intend to have a certain effect on a specific hearer or hearers, then language might already have been shown to be social to the extent of requiring the existence of at least two people (since it is arguable that one could not intend to have an effect on a specific other person unless such a person existed). i shall not take this direct and tempting line here. nevertheless, we are in a position to say that if communication succeeds, there must be these intentions on the part of the speaker, and therefore *if* successful communication is essential to meaning, these intentions are essential to meaning, the presence of intentions is important, since it gives content to an attribution of error by allowing for the possibility of a discrepancy between intention and accomplishment. intention, like belief and expectation, does not require attention or reflection, and intentions are not usually arrived at by conscious reasoning. intentions are not normally attended by any special feelings, nor is our knowledge of our own intentions arrived at (usually) by inference or resort to observation. yet intention has an indefinitely large scope, for intentions depend on the belief that one can do what one intends, and this requires that one believe nothing will prevent the intended action. thus intention would seem to have just the properties needed to make sense of the idea that a speaker has failed to go on as before.6

the view i have just sketched deals only with interpretation, and so presupposes a social environment rather than providing an argu-

<sup>&</sup>lt;sup>s</sup> this is clearly inadequate as it stands. it can he improved by adding the gricean condition that the speaker intends the interpreter to arrive at the right truth conditions through the interpreter's recognition of the speaker's intention to be so interpreted. i do not argue here for the assumption that knowledge of truth conditions is adequate for interpretation.

<sup>&</sup>lt;sup>6</sup> essentially these points about intention are made by crispin wright in attempting, like me, to defuse kripke's view that he has extracted an essentially insoluble 'skeptical paradox' from wittgenstein's treatment of meaning. see crispin wright, 'kripke's account of the argument against private language'.

ment for it. nevertheless, it will be useful at this point to consider certain aspects of the view i think kripke attributes to wittgenstein. (for expository purposes i shall call this kripke's view. in fact kripke does not clearly say he endorses it; and i am uncertain that it is wittgenstein's view. so perhaps it is no one's view.) kripke concentrates on the idea of following a rule. according to this idea, to speak a language is to follow rules. the rules specify what it is to go on 'in the same way'; how, for example, to use a word. there is, however, no inner mental act or process of 'grasping' or of 'following' the rule, so no study or knowledge of what is inside the speaker will reveal whether she is following one set of rules or another. interpreters simply judge that a speaker is following the same rule they (her interpreters) are if the speaker goes on as they would. put in terms of meaning: we judge that a speaker means what we would if we were to utter the same words if she goes on as we would.<sup>7</sup>

we ought to question the appropriateness of the ordinary concept of following a rule for describing what is involved in speaking a language. when we talk of rules of language, we usually have in mind grammarians' or linguists' descriptions (generalized and idealized) of actual practice, or prescriptions grammarians wish we would follow. rules can be a help in learning a language, but their aid is available only in the acquisition of a second language. most learning of how to use words is accomplished without explicitly learning any rules at all.<sup>8</sup> wittgenstein does, of course, treat meaning something in much the same way he does following some procedure, such as adding in arithmetic. but there is a clear distinction between the cases, which explains why we normally use the word 'rule' in one

<sup>8</sup> it should be obvious that the claim that there are internalized or genetically implanted rules of *grammar* is irrelevant here; wittgenstein's and kripke's 'rules' concern what the words of particular languages mean.

<sup>&</sup>lt;sup>7</sup> '... what do i mean when i say that the teacher judges that, for certain cases, the pupil must give the "right" answer? i mean that the teacher judges that the child has given the same answer that he himself would give ... if, in enough concrete cases. jones's incli nations agree with smith's, smith will judge that jones is indeed following the rule' (kripke, *wittgenstein on rules and private language*, 90-1). the following from wittgenstein may bear out this interpretation: 'a person goes by [is guided by] a sign-post only in so far as there exists a regular use of sign-posts, a custom... is what we call "obeying a rule" something that it would be possible for only *one* man to do, and to do only *once* in his life?...—to obey a rule, to make a report, to give an order, to play a game of chess, are *customs* (uses, institutions)' *(philosophical investigations*, §§198, 199). I have ignored a very important aspect of kripke's discussion, his claim that wittgenstein's 'solution' to the problem of meaning is 'skeptical'.

case and not in the other. in the case of adding, there is an explicit procedure for arriving at an answer; we can learn and describe the procedure, and it is appropriate to call the procedure or its description a rule. we normally follow no procedure in speaking; nothing in normal speaking corresponds to adding a column of numbers. if the concept of following a rule is not quite appropriate to describe meaning something by saying something, it is also questionable whether, even if we agree that the use of a language requires a social setting, we should accept without question the idea that meaning something demands (as opposed to sometimes involving) a convention, custom, or institution.<sup>9</sup>

a more important question concerns the idea that linguistic communication requires that a speaker go on in the same way as others do-that to mean something in speaking, one must mean the same thing by the same words as others do. the account i gave above of the sort of expectations that must be satisfied if one person is to understand another did not suggest that those people would have to speak the same language. nor is it clear why this is necessary. perhaps language would never have come into existence unless it could depend on the natural tendency of animals to imitate each other. this may be so, though i have my doubts, but surely it could have been otherwise. if you and i were the only speakers in the world, and you spoke sherpa while i spoke english, we could understand one another, though each of us followed different 'rules' (regularities). what would matter, of course, is that we should each provide the other with something understandable as a language. this is an intention speakers must have; but carrying out this intention, while it may require a degree of what the other perceives as consistency, does not involve following shared rules or conventions. it might even be that because of differences in our vocal chords we couldn't make the same sounds, and therefore couldn't speak the same language. i know of no argument that shows that under such circumstances communication could not take place. so, while it may be true that speaking a language requires that there be an interpreter, it doesn't follow that more than one person must speak the same language. this is fortunate, since if we are precise about what consti-

<sup>&</sup>lt;sup>9</sup> i have expressed my skepticism about the explanatory power of the concepts of rulefollowing and convention in the study of language at greater length in 'communication and convention', essay 18 in *inquiries into truth and interpretation*.

lutes a language, it is probably the case that no two people actually do speak the same language. i conclude that kripke's criterion for speaking a language cannot be right; speaking a language cannot depend on speaking as someone else does (or as many others do).<sup>10</sup>

let us suppose, then, that the test for speaking a language is modified to accommodate this point: speaking a language, we will now claim, does not depend on two or more speakers speaking in the same way; it merely requires that each speaker intentionally make himself interpretable to the other (the speaker must 'go on' more or less as the other expects, or at least is equipped to interpret).

this is certainly a necessary condition for successful *communication*. but why is it a condition that must be satisfied in order to speak a language at all? why couldn't someone go on in the same way satisfy all the conditions for being *interpretable*—without actually being interpreted?"

it is true that our evidence that someone speaks a particular language is mainly based on the fact that he goes on as we expect a speaker of that language to go on. there are, of course, many other sorts of evidence; the speaker's clothes, his companions, his location on the face of the earth, may all be clues to his language. still, we can agree that pinning matters down must in the end depend on the details of speech behavior. the trouble is that the original question concerned neither the conditions for communication nor the question what evidence one person could have that another was speaking a particular language; the issue was why a speaker's sole or first language could not be private.

i have argued in 'a nice derangement of epitaphs' that communication does not demand that languages be shared. the same position is endorsed by noam chomsky, *language and problems of knowledge*. 36-7. for a more extended treatment of this subject directly aimed at kripke. see chomsky. *knowledge of language: its nature, origin and use*, 223-37.

Il kripke seems to allow that robinson crusoe might be judged to be speaking a language, as long as he *could* be included in a society, even if he is not in fact (ever?) in a social setting, there would still have to be a society (or at least another person) legiti mately to judge that robinson crusoe meant something by his noises. chomsky thinks that by allowing the robinson crusoe case kripke contradicts his main thesis. perhaps so: but i think chomsky is wrong in thinking the pure robinson crusoe case possible. by the pure case, i mean a robinson crusoe who has never been in communication with others.

<sup>10</sup> my mother kept a record, as fond mothers will, of the language i spoke when i was 3 years old. it certainly was not the language spoken by others in the family—or by anyone else, i imagine. but she claims, rightly, that she understood me and i her. and it is common for people who cannot or will not speak or write french to answer letters written in french in english.

#### 116 *intersubjective* <-.

our discussion has led to a modification or elucidation of the concept of a private language: i am taking this to mean, not a language only one person speaks, but a language only one person understands. the question now is, why can't there be a language only one person understands?<sup>12</sup> the answer wittgenstein seems to offer in the passage at the head of this essay is: without an interpreter no substance can be given to the claim that the speaker has gone wrong—that he has failed to go on in the same way.

but haven't we, by eliminating the condition that the speaker must go on as the interpreter (or others) would, at the same time inadvertently destroyed all chance of characterizing linguistic error? if there is no social practice with which to compare the speaker's performance, won't whatever the speaker says be, as wittgenstein remarks, in accord with some rule (i.e. in accord with some language)? if the speech behavior of others doesn't provide the norm for the speaker, what can? the answer is that the intention of the speaker to be interpreted in a certain way provides the 'norm'; the speaker falls short of his intention if he fails to speak in such a way as to be understood as he intended. under usual circumstances a speaker knows he is most apt to be understood if he speaks as his listeners would, and so he will intend to speak as he thinks they would. he will then fail in one of his intentions if he does not speak as others do. this simple fact helps explain, i think, why many philosophers have tied the meaning of a speaker's utterances to what others mean by the same words (whether 'others' refers to a linguistic community, experts, or an elite of one sort or another<sup>13</sup>). on my account, this tie is neither essential nor direct; it comes into play only when the speaker intends to be interpreted as (certain) others would be. when this intention is absent, the correct understanding of a speaker is unaffected by usage beyond the intended reach of his voice. (a failed intention to speak 'correctly', unless it foils the intention to be interpreted in a certain way, does not matter to what the speaker means.<sup>14</sup>)

what these considerations show, if they are right, is that there is a

 $^{\mathbb{R}}$  of course there can be a 'language' only one person understands, for example a secret code used in a diary. the question is whether a first language can be private.

<sup>13</sup> for examples, see hilary putnam, the meaning of "meaning" '; tyler burge, 'individualism and the mental': michael dummett, 'the social character of meaning', in *truth and other enigmas*, and ' "a nice derangement of epitaphs": some comments on davidson and hacking', in ernest lepore (ed.), *truth and interpretation: perspectives on the philosophy of donald davidson.* 

this issue is also discussed in essay 2.

;••'.,••. , , , .

weaker and more plausible alternative to kripke's proposed account of what is required in order to mean something by what one says. for while kripke's account makes the test of whether a speaker means something depend on his doing what others do, the same distinction between thinking one means something and actually meaning it can be made in terms of the success of the speaker's intention to be interpreted in a certain way. both ways of making the distinction depend on a social setting, but the second makes different demands on the speaker.<sup>15</sup>

have we now shown that there cannot be a private language? surely not. if we assume that kripke's proposal is correct, then it is true that one way of distinguishing between thinking one means something and actually meaning it requires that language be public; the same can be said for the alternative i have suggested. but nothing definitive has been said to show there may not be some other way of drawing the distinction, a way that does not depend on a social environment.

if we are to establish the essentially public character of language, we need an entirely different sort of argument. in the remainder of this essay i suggest such an argument, an argument that applies not only to speech but also to belief, intention, and the rest of the propositional attitudes. the argument that follows does not start with a skeptical doubt to which an answer is sought, but it does end with what may be wittgenstein's conclusion: language is necessarily a social affair.

consider first a primitive learning situation. some creature is taught, or anyway learns, to respond in a specific way to a stimulus or a class of stimuli. the dog hears a bell and is fed; presently it salivates when it hears the bell. the child babbles, and when it produces a sound like 'table' in the evident presence of a table, it is rewarded; the process is repeated and presently the child says 'table' in the presence of tables. the phenomenon of generalization, of perceived similarity, plays an essential role in the process. one ring of the bell is enough like another to provoke similar responses in the dog, just as one

<sup>&</sup>lt;sup>15</sup> there is a point here that i have not accommodated. a speaker fails in an intention if he is not interpreted as he intends. but it would be wrong to say that such a failure is necessarily a failure to give the meaning to his words that he intended the interpreter to catch. the latter failure depends (in ways that ordinary usage may not definitively settle) on such questions as whether the speaker was justified in believing his interpreter could, or would, interpret him as he intended.

presentation of food is enough like another to engender salivation. if some such discriminative mechanisms were not in our genes, none could be learned. the same goes for the child: we can class the child's stimuli by the similarity of the responses those stimuli elicit in the child.<sup>16</sup>

this seems straightforward, but as psychologists have noticed, there is a problem about the stimulus. in the case of the dog, why say the stimulus is the ringing of the bell? why couldn't it be the vibration of the air close to the ears of the dog—or even the stimulation of its nerve endings? certainly if the air were made to vibrate, in the same way the bell makes it vibrate, it would make no difference to the response of the dog. and if the right nerve endings were activated in the right way, there still would be no difference. in fact, if we must choose, it seems that the proximal cause of the behavior has the best claim to be called the stimulus, since the more distant an event is causally from its perceiver, the more chance there is that the causal chain will be broken. why not say the same about the child: that its responses are not to tables but to patterns of stimulation at its surfaces, since those patterns of stimulation always produce the response, while tables produce it only under favorable conditions?

what explains the fact that it seems so natural to say the dog is responding to the bell, the child to tables? it seems natural to us because it *is* natural—to us. just as the dog and the child respond in similar ways to certain stimuli, so do we. it is we who find it natural to group together the various salivations of the dog; and the events in the world that we effortlessly notice and group together that are causally linked to the dog's behavior are ringings of the bell. we find the child's mournings of 'table' similar, and the objects in the world we naturally class together that accompany those mouthings is a class of tables. the acoustical and visual patterns that speed at their various rates between bell and dog ears, tables and child eyes, we cannot easily observe, and if we could we might have a hard time saying what made them similar. (except by cheating, of course: they are the patterns characteristic of bells ringing, of tables viewed.) nor do we observe the stimulation of nerve endings of other people and animals, and if we did we might find it impossible to describe in a

<sup>&</sup>lt;sup>16</sup> nothing here depends on the amateur psychology involving babbling, differential rewards, or preverbal induction. all that matters is the fact that generalization takes place in one way rather than another, and in similar ways in different people.

noncircular way what made the patterns relevantly similar from trial to trial. the problem would be much the same as the (insoluble) problem of defining tables and bell-ringings in terms of sense data without mentioning tables or bells.

involved in our picture there are now not two but three similarity patterns. the child finds tables similar; we find tables similar; and we find the child's responses in the presence of tables similar. it now makes sense for us to call the responses of the child responses to tables. given these three patterns of response we can assign a location to the stimuli that elicit the child's responses. the relevant stimuli are the objects or events we naturally find similar (tables) which are correlated with responses of the child we find similar. it is a form of triangulation: one line goes from the child in the direction of the table, one line goes from us in the direction of the table, and the third line goes between us and the child. where the lines from child to table and us to table converge, 'the' stimulus is located. given our view of child and world, we can pick out 'the' cause of the child's responses. it is the common cause of our response and the child's response.<sup>17</sup>

enough features are in place to give a meaning to the idea that the stimulus has an objective location in a common space; but nothing in this picture shows that either we, the observers, or our subjects, the dog and the child, have this idea. nevertheless, we have come a good distance. for if i am right, the kind of triangulation i have described, while not sufficient to establish that a creature has a concept of a particular object or kind of object, is necessary if there is to be any answer at all to the question what its concepts are concepts of. if we consider a single creature by itself, its responses, no matter how complex, cannot show that it is reacting to, or thinking about, events a certain distance away rather than, say, on its skin. the solipsist's world can be any size; which is to say, from the solipsist's point of view it has no size, it is not a world.

the problem is not, i should stress, one of verifying what objects or events a creature is responding to; the point is that without a second creature responding to the first, there can be no answer to the question. and of course if there is no answer to this question, there

<sup>&</sup>lt;sup>17</sup> the triangulation metaphor is introduced in essay 7. the idea of the 'common cause' is developed in essay 10. that thought is a social phenomenon is stressed in 'thought and talk', essay 11 in *inquiries into truth and interpretation*, and in essay 7 in this volume.

is no answer to the question what language a creature speaks, since to designate a language as one being spoken requires that utterances be matched up with objects and events in the world (and not, in general, events on the surface of the skin). so we can say, as a preliminary to answering the question with which we began, that before anyone can speak a language, there must be another creature interacting with the speaker. of course this cannot be enough, since mere interaction does not show how the interaction matters to the creatures involved. unless the creatures concerned can be said to react to the interaction, there is no way *they* can take cognitive advantage of the three-way relation which gives content to *our* idea that they are reacting to one thing rather than another.

here is part, i think, of what is required. the interaction must be made available to the interacting creatures. thus the child, learning the word 'table', has already in effect noted that the teacher's responses are similar (rewarding) when its own responses (mouthing 'table') are similar. the teacher on his part is training the child to make similar responses to what he (the teacher) perceives as similar stimuli. for this to work, it is clear that the innate similarity responses of child and teacher—what they naturally group together—must be much alike; otherwise the child will respond to what the teacher takes to be similar stimuli in ways the teacher does not find similar. a condition for being a speaker is that there must be others enough like oneself.

so far i have left out of explicit account the concepts of belief and intention which are clearly essential to speaking a language. i have no thought of trying to introduce these concepts in terms of the simple conditioning situations i have been describing; the concept of thought is not reducible to anything else, much less to these simple concepts. all 1 have tried to show is that interaction among similar creatures is a necessary condition for speaking a language.

now to put two points together. first, if someone is the speaker of a language, there must be another sentient being whose innate similarity responses are sufficiently like his own to provide an answer to the question, what is the stimulus to which the speaker is responding? and second, if the speaker's responses are linguistic, they must be knowingly and intentionally responses to specific stimuli. the speaker must have the concept of the stimulus—of the bell, or of tables. since the bell or a table is identified only by the intersection of two (or more) sets of similarity responses (lines of thought, we might almost say), to have the concept of a table or a bell is to recognize the existence of a triangle, one apex of which is oneself, the second apex another a creature similar to oneself, and the third an object (table or bell) located in a space thus made common.

the only way of knowing that the second apex of the triangle the second creature or person—is reacting to the same object as oneself is to know that the other person has the same object in mind. but then the second person must also know that the first person constitutes an apex of the same triangle another apex of which the second person occupies. for two people to know of each other that they are so related, that their thoughts are so related, requires that they be in communication. each of them must speak to the other and be understood by the other. they don't, as i said, have to mean the same thing by the same words, but they must each be an interpreter of the other.

the hasty remarks of the last two paragraphs indicate the sort of work that would be necessary to give an account of meaningful speech. but such an account was not my aim; i was looking only to find an argument why a first language cannot be private.

the argument shows that one's first language cannot be a private language, that is, a language understood by only one creature, and to this extent it is in agreement with kripke's wittgenstein. but the argument takes a different course, and the flavor it gives the social aspect of language is different. kripke depends on the second person, or a community, to embody a routine which the speaker can share. in contrast, the argument i have outlined does not require (though of course it allows) a shared routine, but it does depend on the interaction of at least two speaker-interpreters, for if i am right, there would be no saying what a speaker was talking or thinking about, no basis for claiming he could locate objects in an objective space and time, without interaction with a second person.

the considerations i have put forward do not apply to language only; they apply equally to thought in general. belief, intention, and the other propositional attitudes are all social in that they are states a creature cannot be in without having the concept of intersubjective truth, and this is a concept one cannot have without sharing, and knowing that one shares, a world, and a way of thinking about the world, with someone else.

# 9 the emergence of thought

according to the bible, the word was there from the beginning, but it seems more plausible that words and thought emerged rather recently. what does it mean, though, to say some aspect of the world 'emerged'?

emergence makes sense only as seen from a point of view, from within a set of concepts. given our set of concepts, we can appreciate the idea that different concepts were instantiated at different times. the concepts of quantum physics were instantiated very early; the concepts of the various elements were instantiated over a longer subsequent period. the concept of life was instantiated quite recently, and the concepts of thought and language were filled out only moments ago, with the coming of the higher mammals. this doesn't mean the laws of physics changed, but it does mean that in order to describe and explain thought we need concepts that can't be defined in the vocabulary of physics (or any of the natural sciences). thus, in a clear sense, emergence is relative to a set of concepts. concepts themselves are abstractions and so timeless, of course, but it can happen that certain concepts are instantiated only as time goes on. the universe is, needless to say, indifferent to our concepts. but we care very much about the instantiation of the concepts on which the concept of thought depends: such concepts as belief, desire, intention, intentional action, memory, perception, and all the rest of our human attitudes and attributes.

let me begin by saying why it is that in my view to try to say anything interesting and deep about the emergence of thought is so difficult. the reason is that there are so many concepts that we must have in order to talk about or describe thinking, acting on a reason, believing, or doubting, all of which depend on each other. this is the

*holism of the mental,* the interdependence of various aspects of mentality. within any one dimension of mentality, such as belief, it seems clear that it is impossible to take an atomistic approach, because it is impossible to make sense of the idea of having only one or two beliefs. beliefs do not come one at a time: what identifies a belief and makes it the belief that it is is the relationship (among other factors) to other beliefs. one cannot believe that he or she is seeing a cat without believing many other things: one must know what a cat is, what seeing is, and above all, one must recognize the possibility, however remote, that one may be wrong. some people suppose that a dog might have such an isolated belief, but it seems to me, for the very reasons i am now rehearsing, that dogs do not have beliefs, or any other propositional attitudes. they do not form judgements.

the reason neither a dog nor any other creature can have a single belief, such as that it is seeing a cat, is that what identifies a belief is what we loosely call its *propositional content*. thus, to have a belief about a cat, one must have mastery of the concepts that are involved in this judgement or belief. a creature does not have the concept of a cat merely because it can discriminate cats from other things in its environment. mice are very good at telling cats apart from trees, lions, and snakes. but being able to discriminate cats is not the same thing as having the concept of a cat. you have the concept of a cat only if you can make sense of the idea of misapplying the concept, of believing or judging that something is a cat which is not a cat. to have the concept of a cat, you must have the concept of an animal, or at least of a continuing physical object, the concept of an object that moves in certain ways, something that can move freely in its environment, something that has sensations, there is no fixed list of things you have to know about, or associate with, felinity; but unless you have a lot of beliefs about what a cat is, you don't have the concept of a cat.

because of the fact that beliefs are individuated and identified by their relations to other beliefs, one must have a large number of beliefs if one is to have any. beliefs support one another, and give each other content. beliefs also have logical relations to one another. as a result, unless one's beliefs are roughly consistent with each other, there is no identifying the contents of beliefs. a degree of rationality or consistency is therefore a condition for having beliefs.

this is not to say everyone is perfectly rational; anyone is capable of making a mistake in logic, or of entertaining beliefs that are inconsistent with each other. (both frege and quine, for example, wrote books based on inconsistent logics.) but there is a limit to how inconsistent a person can be and still be credited with clearly defined attitudes. each belief must be involved with other beliefs with which it is consistent if it is to be identified as having a clear content. inconsistency, or other forms of irrationality, can occur only within the space of reasons; inconsistencies are perturbations of rationality, not mere absence of rationality.

i am insisting, then, that one must have a quite fully developed set of basic concepts in order to have any concepts at all. and i have hardly begun to sketch the degree of interdependence among beliefs. for there are, in addition to the logical relations, relations of evidential support. while few concepts can be defined in terms of their empirical content, belief in some factual matters must count in favor of the truth of others; such relations of evidential or inductive support are essential to the identity of many beliefs.

in addition to the relations among beliefs, there are the relations between beliefs on the one hand and evaluative attitudes on the other. it is doubtful that a creature could be credited with beliefs if it did not also have desires, for it is an essential aspect of beliefs that they affect, and are evidenced by, behavior. evaluative attitudes, including desires, intentions, moral convictions, views about duties, and obligations, are also propositional attitudes, this is not always obvious from the grammar of the expressions we use to attribute such attitudes. we speak, for example, of wanting an apple. if someone says, 'i want an apple', it may be rude but it would always be appropriate to ask, 'and what do you want it for?' it makes no sense to say, 'i just want the apple, not for anything or in any particular way'. i might reply, 'all right, you've got it. there it is in the tree. what more do you want?' the answer could be, 'i want it in my hand; or in my basket; or detached from the tree and in my stomach, etc.'. wants and desires are directed to propositional contents. what one wants is *that* it be true that one has the apple in hand, or ready to eat. similarly for intentions. someone who intends to go to the opera intends to make it the case *that* he is at the opera.

beliefs and desires conspire to cause, rationalize, and explain intentional actions. we act intentionally for reasons, and our reasons always include both values and beliefs. we would not act unless there were some value or end we hoped to achieve (or some supposed evil we hoped to avoid), and we believed our course of action was a way of realizing our aim. decision theory is a way of

systematizing the relations among beliefs, desires, and actions. it does this by imposing a complex, but clearly defined, pattern on the way in which people's beliefs and desires interact. decision theory is often derided as a false description of how people actually act; great ingenuity has gone into tricking subjects into making a series of choices inconsistent according to the theory, as a description of actual behavior, decision theory seems false because there is no completely satisfactory way to test it; the test always depends on exactly how the theory is given empirical application. nevertheless, the theory answers to our intuitions about how actual decisions are made; in effect it simply spells out our commonsense apparatus for explaining intentional action. for we all, whether we think about it or not, make our decisions in terms of how we weigh the values of various possible outcomes of our actions, and how likely we think one or another course of action is to attain those values. we understand why someone acts as he does only by supposing that he or she values to various degrees the possible results of action, and how probable he or she thinks a given action is to produce one or another result. so although not capable of precise predictions of actual choices, decision theory nevertheless corresponds to our intuitions about how actual decisions are made, and so is part of our commonsense apparatus for explaining intentional behavior.

in these remarks, i am emphasizing the holism of the mental, the extent to which various aspects of the mental depend on each other. there are, as i have argued, no beliefs without many related beliefs, no beliefs without desires, no desires without beliefs, no intentions without both beliefs and desires. conceptually, actions themselves belong to the realm of the mental, for a piece of behavior counts as an action only if there is some description under which it is intentional, and so can be explained as done for a reason. of course, any action will be unintentional under many descriptions. like any object or event, actions can be described in endless ways, thus if i intentionally sit down, i unintentionally move some air molecules out of the way; i unintentionally depress the cushion on which i sit; i unintentionally change the center of gravity of the earth, the solar system, and our home galaxy, even our failures are intentional, not that we intend to fail, of course; the point is that in failing there nevertheless is something we are doing on purpose. if my return of serve goes long, an intention failed: still, i did intentionally strike the ball

# the emergence of thought 127

you will see my strategy. i am enlarging the field of the mental aspects of life all of which must emerge together. intentional action cannot emerge before belief and desire, for an intentional action is one explained by beliefs and desires that caused it; beliefs can't emerge one at a time, since the content of each belief depends on its place in the nexus of further beliefs; and so on. it is the holism of the mental that makes its emergence so difficult to describe. there cannot be a sequence of emerging features of the mental, not if those features are to be described in the usual mentalistic vocabulary. of course, everything in the universe and its history can in principle be described in the language of physics, and so each stage in the emergence of thought can be described in physical terms. but this will fail as an explanation of the emergence of the mental since we do not have, and cannot expect to find, a way of mapping events described in the physical vocabulary onto events described in the mental vocabulary.

the difficulty in describing the emergence of mental phenomena is a conceptual problem: it is the difficulty of describing the early stages in the maturing of reason, the stages that precede the situation in which concepts like intention, belief, and desire have clear application. in both the evolution of thought in the history of mankind, and the evolution of thought in an individual, there is a stage at which there is no thought followed by a subsequent stage at which there is thought. to describe the emergence of thought would be to describe the process which leads from the first to the second of these stages. what we lack is a satisfactory vocabulary for describing the intermediate steps. we are able to describe what a preverbal child does by employing the language of neurology, or in crude behavioristic terms we can describe movements and the sounds emitted. you can deceive yourself into thinking that the child is talking if it makes sounds which, if made by a genuine language user, would have a definite meaning. (it is even possible to do this with chimpanzees.) but words, like thoughts, have a familiar meaning, a propositional content, only if they occur in a rich context, for such a context is required to give the words or thoughts a location and a meaningful function. if a mouse had vocal cords of the right sort, you could train it to say 'cheese'. but that word would not have a meaning when uttered by the mouse, nor would the mouse understand what it 'said'. infants utter words in this way; if they did not, they would never come to have a language. but if you want to describe what is going

on in the head of the child when it has a few words which it utters in appropriate situations, you will fail for lack of the right sort of words of your own. we have many vocabularies for describing nature when we regard it as mindless, and we have a mentalistic vocabulary for describing thought and intentional action; what we lack is a way of describing what is in between. this is particularly evident when we speak of the 'intentions' and 'desires' of simple animals. we have no better way to explain what they do.

it is not that we have a clear idea what sort of language we could use to describe half-formed minds; there may be a very deep conceptual difficulty or impossibility involved. that means there is a perhaps insuperable problem in giving a full description of the emergence of thought. i am thankful that i am not in the field of developmental psychology.

despite these pessimistic remarks, i do have some suggestions about how we might approach the problem of saying something intelligible about the emergence of thought, there is a prelinguistic, precognitive situation which seems to me to constitute a necessary condition for thought and language, a condition that can exist independent of thought, and can therefore precede it. both in the case of nonhuman animals and in the case of small children, it is a condition that can be observed to obtain. the basic situation is one that involves two or more creatures simultaneously in interaction with each other and with the world they share; it is what i call triangulation. it is the result of a threefold interaction, an interaction which is twofold from the point of view of each of the two agents: each is interacting simultaneously with the world and with the other agent. to put this in a slightly different way, each creature learns to correlate the reactions of other creatures with changes or objects in the world to which it also reacts. one sees this in its simplest form in a school of fish, where each fish reacts almost instantaneously to the motions of the others. this is apparently a reaction that is wired in. a learned reaction can be observed in certain monkeys which make three distinguishable sounds depending on whether they see a snake, an eagle, or a lion approaching; the other monkeys, perhaps without seeing the threat themselves, react to the warning sounds in ways appropriate to the different dangers, by climbing trees, running, or hiding, but on reflection we realize that the behavior of these primates, complex and purposeful as it is, cannot be due to propositional beliefs, desires, or intentions, nor does their mode of communication constitute a language.

#### the emergence of thought 129

nevertheless, the triangle i have indicated is essential to the existence, and hence to the emergence, of thought. for without the triangle, there are two aspects of thought for which we cannot account. these two aspects are the objectivity of thought and the empirical content of thoughts about the external world.

thought, propositional thought, is objective in the sense that it has a content which is true or false independent (with rare exceptions) of the existence of the thought or the thinker. furthermore, this is a fact of which a thinker must be aware; one cannot believe something, or doubt it, without knowing that what one believes or doubts may be either true or false and that one may be wrong, where do we get the idea that we may be mistaken, that things may not be as we think they are? wittgenstein has suggested, or at least i take him to have suggested, that we would not have the concept of getting things wrong or right if it were not for our interactions with other people. the triangle i have described stands for the simplest interpersonal situation. in it two (or more) creatures each correlate their own reactions to external phenomena with the reactions of the other. once these correlations are set up, each creature is in a position to expect the external phenomenon when it perceives the associated reaction of the other, what introduces the possibility of error is the occasional failure of the expectation: the reactions do not correlate, wittgenstein expresses this idea when he talks of the difference between following a rule and merely thinking one is following a rule; he says that following the rule (getting things right) is at bottom a matter of doing as others do.' of course, the others may sometimes be wrong. the point isn't that consensus defines the concept of truth but that it creates the space for its application. if this is right, then thought as well as language is necessarily social.

social interaction, triangulation, also gives us the only account of how experience gives a specific content to our thoughts. without other people with whom to share responses to a mutual environment, there is no answer to the question what it is in the world to which we are responding. the reason has to do with the ambiguity of the concept of cause. it is essential to resolve these ambiguities, since it is, in the simplest cases, what causes a belief that gives it its content. in the present case, the cause is doubly indeterminate: with respect to width, and with respect to distance. the first ambiguity concerns

<sup>1</sup> 'but see the caveats of essay 8.

how much of the total cause of a belief is relevant to content. the brief answer is that it is the part or aspect of the total cause that typically causes relevantly similar responses. what makes the responses relevantly similar in turn is the fact that others find those responses similar; once more it is the social sharing of reactions that makes the objectivity of content available. the second problem has to do with the ambiguity of the relevant stimulus, whether it is proximal (at the skin, say) or distal. what makes the distal stimulus the relevant determiner of content is again its social character; it is the cause that is shared. the stimulus is thus triangulated; it is where causes converge in the world.

the triangular relationship between agents and an environment to which they mutually react is, i have argued, necessary to thought. it is not sufficient, as is shown by the fact that it can exist in animals we do not credit with judgement. for this reason we are in a position to say something about a situation that must exist if thought does, but it is a situation that can exist independently, and so can precede thought in the order of things. it can exist first, and it surely does. thus we can say that a certain kind of primitive social interaction is part of the story of how thought emerged.

what more is needed for thought? i think the answer is language. in itself, this is not much help, since it is obvious that a creature that has a language can think; language is an instrument for the expression of propositional contents. still, we can ask why language is essential to thought. the reason, stated briefly, is that unless the base line of the triangle, the line between the two agents, is strengthened to the point where it can implement the communication of propositional contents, there is no way the agents can make use of the triangular situation to form judgements about the world. only when language is in place can creatures appreciate the concept of objective truth. there is much more that can be said about the deep relationship between language as a means of interpersonal communication and thought, but i shall here assume this relationship, and make use of it to say something more about the emergence of thought. language offers us the opportunity to make relatively sharp comparisons among various linguistic and prelinguistic symbol systems, and various ways of describing them.

for this purpose, i suggest that we think of formal semantic theories in analogy with theories of fundamental measurement. in a theory of fundamental measurement, there are one or more primitive concepts; in the case of the measurement of weight, for example, these may be the relation of one thing being *at least us heavy as* another and the operation of *adding* one thing to another (this is a function that maps two things on to their sum). these concepts are not defined, but before the theory can be applied to, or tested against, some collection of objects, a method for giving an empirical content to the primitives must be indicated. (thus *a* is at least as heavy as *b if a* does not go up in the pan of an equal-arm balance which holds *b* in the other pan.) a set *of axioms* then specifies the logical properties of the primitives and the entities to be measured. (so the axioms for weight will specify that the relation of being at least as heavy as the sum of *b* and *a*.)

the aim of the theory is to describe in precise terms the sort of structure a set of entities must have if we are to assign weights to them. if weights are to have the properties we intuitively expect, the axioms must be proven sufficient to account for these properties. one of the things we want to prove *(about* the theory, not *in* the theory) is a representation theorem, which states that numbers can be assigned to the objects covered by the theory in a way that keeps track of the relations in weight among the different objects. the second thing to be proven is a uniqueness theorem, which states that if one set of numbers represents the weights of objects, then any other set of numbers that is related to the first set by a certain kind of transformation will represent the weights just as well. in the case of weight, multiplication by any positive constant will transform one measure into another; weights in pounds are as good as weights in kilos.

now let me compare a semantic theory for some speaker or group of speakers. tarski's truth definitions, when modified to apply to natural languages, may be viewed as such theories. like a theory for the measurement of weight, a semantic theory will place restrictions on one or more primitive concepts. the most important primitive is that of truth. (here i deviate from tarski's interests and intentions. where he *defined* truth for particular languages, i have in mind axiomatizations of a general, and hence indefinable, concept of truth. this course risks what tarski sought to avoid, namely inconsistency.) the theory imposes a certain structure on any language it can describe; in giving the truth conditions of all the sentences of a language, it necessarily defines the logical relations among sentences.

we can think of such a theory as saying what is needed in order to understand a language; it describes what a speaker intuitively knows about his own language, and it can serve to interpret what such a speaker says. the adequacy of the theory for these purposes can be shown by demonstrating that one can represent the properties of the object language in some language one understands. thus the theory serves to validate a translation manual. just as we use numbers as a convenient way of keeping track of the weights of objects, we keep track of the semantic features of a language by using sentences we understand. what makes the numbers so suited to this job is that we know exactly what structure they have. in much the same way, the sentences of a language we understand have a known structure which we can use to understand other speakers.

different forms of measurement make use of different properties of the numbers; roughly speaking, the more properties of the numbers they use, the more powerful or informative the mode of measurement. for example, when numbers are used to keep track of the members of a team, their only significance is that one number is different from another. the mohs hardness scale makes use of the fact that one number is larger than another: any other ten numbers (besides those actually used, 1-10), as long as ordered in ascending magnitude, would do as well. any linear transformation of the fahrenheit temperature scale would serve, while the more powerful weight and distance scales allow only multiplication by a positive constant. in the case of temperature, two numbers are arbitrary, the zero point, and the numerical value of an interval; as a result, it makes no sense to say one body is twice as warm as another. on the other hand, one body can weigh twice as much as another.

this gives us a way of thinking about the relative expressive powers of languages. it is intuitively clear that many sentences are equivalent in what they convey. thus it makes sense to ask: in how many different ways could we assign our own sentences to the sentences of another language and still capture the same information (that is, how many of these ways would constitute equally satisfactory translation manuals?). of course, a translation manual doesn't translate one sentence at a time; the manual is a function that maps all of the sentences of one language onto the sentences of another. (just as you can't compare weights in pounds with weights in kilos, so you can't change translation manuals without notice, for this would destroy the relations among sentences.) what we must
consider, then, are total systems of translation. the claim that there are such empirically equivalent systems is quine's thesis of the indeterminacy of translation: you could keep track of everything you could ever learn about what someone means by what they say, and therefore everything that they mean, in different ways.

what is clear is that the more different ways there are of representing another language in our own, the less expressive power that language has (from our point of view, of course). let us consider some examples. all languages have a finite basic vocabulary. we can imagine a language with just proper names and predicates. such a language must have a finite number of sentences, so our theory need do no more than match some sentence of our own to each of the sentences in the other language. the semantics of such a language are simple indeed: there is no need to break sentences down into parts, no reason to employ recursive devices, no grounds for attributing an ontology. if you consider how simple the semantics of such a language can be, you must wonder whether it ought to be called a language; after all, a 'language' that does not require that we find a relation between its parts ('words') and entities in the world barely merits the term.

we advance a step in complexity by introducing the truth-functional connectives like the signs of conjunction, negation, alternation, and the material conditional ('and', 'not', 'or', 'if. . . then . . .'). there now will be an infinity of sentences, since one can add negation to any sentence and get a sentence, or conjoin a sentence to any sentence and produce another sentence. this introduces the element of productivity, the possibility of producing and understanding new sentences on the basis of simple syntactic and semantic rules. the semantics remain elementary, however, since the semantics for our first, finite language still suffice to give the truth conditions of the atomic sentences, and the truth tables give the semantics of the molecular sentences. the only parts of speech that need to be distinguished are the simplest sentences and sentential connectives; ontology is still waiting in the wings.

the last stage requires a leap; it introduces quantification, the concepts expressed by the words 'some' and 'all'. once we advance to this stage, we have arrived at languages that match, or begin to match, our own in complexity. the semantics must distinguish predicates of various orders, singular terms, and expressions that do the work of variables (pronouns and the like), and, of course, the quantifiers. for

# 134 intersubjective

the first time, it is necessary to find structure in the simplest sentences, and to correlate predicates and singular terms with objects. it is here, in my opinion, that we reach the degree of expressive sophistication that we associate with thought, for it is only at this level that there is positive evidence that the speaker of the language can predicate properties of objects and events. it was alfred tarski who made this point clear, for it was he who first showed what was required if we are to characterize truth for languages with general quantification. he did this by introducing the concept *of satisfaction* which is a relation between predicates of any degree of complexity and objects in the world.<sup>2</sup>

it is not necessary to be dogmatic about where thought, or language, begins in order to appreciate the relative clarity that formal semantics and the analogy with measurement theory bring to our thinking about the emergence of thought. for however troubled and vague our attempts to describe the stages leading up to thought and language may be, semantics and measurement theory give us a platform from which to distinguish levels of development from an objective standpoint.

<sup>2</sup> alfred tarski, 'the concept of truth in formalized languages'.

# 10 a coherence theory of **truth** and knowledge

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

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

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

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

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

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

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

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

this way of stating the position can at best be taken as a hint, since there is no useful way to count beliefs, and so no clear meaning to the idea that most of a person's beliefs are true. a somewhat better way to put the point is to say there is a presumption in favor of the truth of a belief that coheres with a significant mass of belief. every belief in a coherent total set of beliefs is justified in the light of this presumption, much as every intentional action taken by a rational agent (one whose choices, beliefs, and desires cohere in the sense of bayesian decision theory) is justified. so to repeat, if knowledge is justified true belief, then it would seem that all the true beliefs of a consistent believer constitute knowledge. this conclusion, though too vague and hasty to be right, contains an important core of truth, as i shall argue. meanwhile i merely note the many problems asking for treatment: what exactly does coherence demand? how much of inductive practice should be included, how much of the true theory (if there is one) of evidential support must be in there? since no person has a completely consistent body of convictions, coherence with *which* beliefs creates a presumption of truth? some of these problems will be put in better perspective presently.

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

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

<sup>&</sup>lt;sup>1</sup> see my 'true to the facts', essay 3 in *inquiries into truth and interpretation*, and 'afterthoughts' in this volume.

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

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

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

<sup>&</sup>lt;sup>2</sup> hilary putnam, meaning and the moral sciences, 125.

<sup>3&#</sup>x27; see my 'on the very idea of a conceptual scheme', essay 13 in *inquiries into truth and interpretation*.

#### a coherence theory 141

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

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

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

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

<sup>&</sup>lt;sup>4</sup> richard rorty, *philosophy and the mirror of nature*. 178.

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

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

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

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

let us try a bolder tack. suppose we say that sensations themselves, verbalized or not, justify certain beliefs that go beyond what is given in sensation. so, under certain conditions, having the sensation of seeing a green light flashing may justify the belief that a green light is flashing. the problem is to see how the sensation justifies the belief. of course if someone has the sensation of seeing a green light flashing, it is likely, under certain circumstances, that a green light is flashing. we can say this, since we know of his sensation, but he can't say it, since we are supposing he is justified without having to depend on believing he has the sensation. suppose he believed he didn't have the sensation. would the sensation still justify him in the belief in an objective flashing green light? the relation between a sensation and a belief cannot be logical, since sensations are not beliefs or other propositional attitudes. what then is the relation? the answer is, i think, obvious: the relation is causal. sensations cause some beliefs and in *this* sense are the basis or ground of those beliefs. but a causal explanation of a belief does not show how or why the belief is justified.

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

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

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

<sup>&</sup>lt;sup>5</sup> w. v. quine, the nature of natural knowledge', 68. many other passages in quine suggest that he hopes to assimilate sensory causes to evidence. in *word and object*, 22, he writes that 'surface irritations . . . exhaust our clues to an external world'. in *ontological relativity*, 75, we find that 'the stimulation of his sensory receptors is all the evidence anybody has had to go on, ultimately, in arriving at his picture of the world.' on the same page: two cardinal tenets of empiricism remain unassailable . . . one is that whatever evidence there *is* for science *is* sensory evidence. the other . . . is that all inculcation of meanings of words, must rest ultimately on sensory evidence.' in *the roots of reference*, 37-8, quine says 'observations' are basic 'both in the support of theory and in the learn ing of language', and then goes on. 'what are observations? they are visual, auditory, tactual, olfactory. they are sensory, evidently, and thus subjective . . . should we say then that the observation is not the sensation . . . '? no . . .' quine goes on to abandon talk of observations sentences, unlike observations, cannot play the role of evidence unless we have reason to believe they are true.

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

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

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

but what concerns me here is that quine and dummett agree on a basic principle, which is that whatever there is to meaning must be traced back somehow to experience, the given, or patterns of sensory stimulation, something intermediate between belief and the usual objects our beliefs are about. once we take this step, we open the door to skepticism, for we must then allow that a very great many perhaps most—of the sentences we hold to be true may in fact be false. it is ironical. trying to make meaning accessible has made truth inaccessible. when meaning goes epistemological in this way, truth and meaning are necessarily divorced. one can, of course, arrange a shotgun wedding by redefining truth as what we are justified in asserting. but this does not marry the original mates.

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

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

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

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

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

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

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

belief can be seen to be veridical by considering what determines the existence and contents of a belief. belief, like the other so-called propositional attitudes, is supervenient on facts of various sorts, behavioral, neurophysiological, biological, and physical. the reason for pointing this out is not to encourage definitional or nomological reduction of psychological phenomena to something more basic, and certainly not to suggest epistemological priorities. the point is rather understanding. we gain one kind of insight into the nature of the propositional attitudes when we relate them systematically to one another and to phenomena on other levels. since the propositional attitudes are deeply interlocked, we cannot learn the nature of one by first winning understanding of another. as interpreters, we work our way into the whole system, depending much on the pattern of interrelationships.

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

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

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

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

the general lines of the solution, like the problem itself, are owed to quine. i will, however, introduce some changes into quine's solution, as 1 have into the statement of the problem. the changes are directly relevant to the issue of epistemological skepticism.

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

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

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

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

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

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

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

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

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

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

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

it is an artifact of the interpreter's correct interpretation of a person's speech and attitudes that there is a large degree of truth and consistency in the thought and speech of an agent. but this is truth and consistency by the interpreter's standards. why couldn't it happen that speaker and interpreter understand one another on the basis of shared but erroneous beliefs? this can, and no doubt often does, happen. but it cannot be the rule. for imagine for a moment an interpreter who is omniscient about the world, and about what does and would cause a speaker to assent to any sentence in his (potentially unlimited) repertoire. the omniscient interpreter, using the same method as the fallible interpreter, finds the fallible speaker largely consistent and correct. by his own standards, of course, but since these are objectively correct, the fallible speaker is seen to be largely correct and consistent by objective standards. we may also, if we want, let the omniscient interpreter turn his attention to the fallible interpreter of the fallible speaker. it turns out that the fallible interpreter can be wrong about some things, but not in general; and so he cannot share universal error with the agent he is interpreting. once we agree to the general method of interpretation i have sketched, it becomes impossible correctly to hold that anyone could be mostly wrong about how things are.

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

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

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

<sup>&</sup>lt;sup>7</sup> it is clear that the causal theory of meaning has little in common with the causal theories of reference of kripke and putnam. those theories look to causal relations between names and objects of which speakers may well be ignorant. the chance of systematic error is thus increased. my causal theory does the reverse by connecting the cause of a belief with its object.

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

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

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

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

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

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

# afterthoughts

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

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

in our 1983 discussion at the pacific division meeting of the american philosophical association i agreed to stop calling my position either a coherence or a correspondence theory if he would give up the pragmatist theory of truth. he has done his part; he now explicitly rejects both james and peirce on truth. i am glad to hold to my side of the bargain. if it had not already been published, i would now change the title of 'a coherence theory of truth and knowledge', and i would not describe the project as showing how 'coherence yields correspondence'. on internal evidence alone, as rorty points out, my view cannot be called a correspondence theory. as long ago as 1969 ('true to the facts'<sup>8</sup>) i argued that nothing can usefully and intelligibly be said to correspond to a sentence; and i repeated this in 'a coherence theory of truth and knowledge'. i thought then the fact that in characterizing truth for a language it is necessary to put words into relation with objects was enough to give some grip for the idea of correspondence; but this now seems to me a mistake, the mistake is in a way only a misnomer, but terminological infelicities have a way of breeding conceptual confusion, and so it is here, correspondence theories have always been conceived as

. in inquiries into truth and interpretation.

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

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

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

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

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

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

<sup>9</sup> hilary putnam, realism and reason, p. xviii.

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

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

# 1 empirical content 1

the dispute between schlick and neurath over the foundations of empirical knowledge illustrates the difficulties in draw trying to epis-temological conclusions from a verificationist theory of meaning. it also shows how assuming the general correctness of science does not automatically avoid, or provide an easy answer to, skepticism. but while neither schlick nor neurath arrived at a satisfactory account of empirical knowledge, there are promising hints of a better theory in their writings. following up these hints, and drawing on further ideas in hempel, carnap, and particularly quine, i suggest the direction i think a naturalistic epistemology should take.

the logical positivists agreed that the empirical content of an interpreted sentence derives from its relations to a subset of sentences that report, or are based on, observation or experience. two main sources of difficulty and dispute immediately became evident. one was the question how to characterize the relations between protocol sentences and other sentences. the history of the developments and changes in the views of the logical positivists and their followers on this problem has been masterfully recorded, as well as much contributed to, by carl hempel.<sup>1</sup> this is not my present subject.

the second question was how protocol sentences should be formulated, and what their relation to experience or observation is. this is the issue i wish to discuss, and on which schlick and neurath disagreed, schlick endorsing a foundationalist epistemology and neurath a coherence theory. the difference was expressed in fairly strong terms. neurath described the foundationalist position as

' see carl hempel, 'empiricist criteria of cognitive significance: problems and changes', in *aspects of scientific explanation*.

'related to the belief in *immediate experiences* which is current in traditional academic philosophy', and remarked that 'methodological solipsism' (carnap's term for a view like schlick's) does 'not become more serviceable because of the addition of the word "methodological" '.<sup>2</sup> schlick in turn called neurath's version of the coherence theory an 'astounding error'.<sup>3</sup>

astounding error or not, carnap and hempel at one time seemed to agree with neurath. in 1935 hempel wrote, 'i think that there is no essential difference left between protocol statements and other statements,' and he concurred with carnap in holding that once the question which sentences were basic was put in the 'formal mode', the answer became a matter for convention to decide. 'this insight', he concluded, echoing carnap's words, 'eliminates from the logical positivist's theory of verification and truth a remainder of absolutism which is due to metaphysical tendencies'.<sup>4</sup> schlick jeered at such conventionalism, saying it made truth as relative as 'all the measuring rods of physics'. in an ironic vein he added, 'and it is this view with its consequences that has been commended as banishing the last remnant of "absolutism" from philosophy'.<sup>5</sup>

it is not entirely clear, however, just where matters stood. fifteen years later, in a quasi-historical article, hempel asserted that the fundamental tenet of modern empiricism is the view that all non-analytic knowledge is based on experience.' he went on to explain what it means for knowledge to be based on experience: nonanalytic knowledge can be expressed by sentences that are confirmed (in a specified way) by observation sentences, which in turn are 'ascertained' to be true by direct observation.<sup>6</sup> this sounds more like schlick, and indeed like the 'fatal confrontation of statements and facts' which hempel had previously rejected.<sup>7</sup>

in a note added to a reprinting of 'studies in the logic of confirmation'<sup>8</sup> hempel suggests a way of partially reconciling the apparently opposed points of view. truth, he and carnap had come to

<sup>3</sup> merit/, schlick, the foundation of knowledge', lp 215.

<sup>4</sup> carl hempel, 'on the logical positivist's theory of truth', 58, 59.

<sup>5</sup> morit/ schlick, the foundation of knowledge', *lp* 213

<sup>6</sup> carl hempel, the empiricist criterion of meaning'. *lp* 108-10.

<sup>8</sup> note 49 to 'studies in the logic of confirmation', in *aspects of scientific explanation*.

<sup>-</sup> otto neurath, 'protocol sentences'; quoted from ayer (ed.). *logical positivism*, 204, 206. all subsequent page references from *logical positivism* are preceded by *lp*.

<sup>&</sup>lt;sup>7</sup> carl hempel, 'on the logical positivist's theory of truth', 51.

realize in the light of tarski's work, is a legitimate semantic notion, and should not be treated as a matter of coherence. confirmation, on the other hand, was of two sorts, *relative* and *absolute*. the logical study of confirmation was the study of the extent to which an arbitrary set of sentences confirmed a hypothesis. in this context, one could only say that relative to a set of sentences (whatever its provenance), a hypothesis was confirmed or disconfirmed. absolute confirmation, on the other hand, depended on a 'pragmatic' decision to treat certain sentences as true. neurath and carnap, hempel suggests, were thinking of relative confirmation, which invites a coherence theory. schlick, and hempel in 1950, were thinking of absolute confirmation.

this important distinction of hempel's is revealing, since it does seem at times that early discussions of protocol sentences vacillated between treating such sentences as any sentences with a specified syntax, and treating them as sentences that were accepted, perhaps on the basis of observation or experience. but the distinction cannot reconcile all the differences. the differences that remained were these: schlick held, while neurath denied, that protocol sentences may be established as true once and for all; schlick claimed, while neurath denied, that a sentence could intelligibly be said to be compared to reality. there were also differences over the question of the proper subject matter of protocols, and the question whether they reported something private or something public. obviously, these various points are closely related to one another.

one way to approach our central problem is to ask what the nature of *evidence* is: does it consist of objects, events, facts, experiences, sensations, beliefs, propositions, or sentences? almost every one of these possible answers can be found in the writings of the vienna circle.

i observe two pieces of green paper [writes schlick] and determine that they have the same color. the proposition which asserts the sameness of color is verified, among other ways, by the fact that at the same time i have two experiences of the same color. the proposition: 'there are two spots of the same color before me now' cannot be reduced to others; it is verified by the fact that it describes the given.<sup>9</sup>

what exactly does the verifying? schlick says it is the fact that he has certain experiences that are veridical. but if the same experiences were not veridical, would they still verify the same proposition? the

<sup>9</sup> moritz schlick, 'positivism and realism'. *lp* 92-3.

fact that i see a piece of paper implies that there is a piece of paper, but here we do not move from evidence to hypothesis in an interesting way; we merely deduce an entailed proposition. elsewhere, schlick insists that we must start with statements that 'have their origin' in observation sentences; and he elaborates this as: 'they derive, as one may confidently say in the traditional way of speaking, "from experience" '.'º in the same essay he declares that 'it is clear, and is so far as i know disputed by no one, that knowledge in life and science in some sense begins with confirmation of facts, and that the "protocol statements" in which this occurs stand in the same sense at the beginning of science." in another passage he says that everything goes back to what is 'immediately observed'.<sup>12</sup> as aver put it, some propositions can be 'directly confronted with the facts'. let me try to bring out in one further way the apparently puzzling question of the ontological status of evidence. we say that laws are confirmed by their positive instances; so the positive instances are, presumably, evidence for the laws they confirm, suppose, for the sake of clarity in one direction, that we identify laws with universally quantified conditional sentences-sentences which are, needless to say, interpreted. what is a positive instance? let the law have the form  $'(x)(fx^*gx)$  then, hempel suggests, it is reasonable to suppose that an *object* that is f and g confirms the law.<sup>13</sup> (goodman often talks this way in fact, fiction, and forecast.) hempel then explains that instead of viewing confirmation as a 'relation between an object or an ordered set of objects, representing the evidence, and a sentence, representing the hypothesis', he will take it to be a relation between a sentence that describes the evidence, and the hypothesis, thus, the evidence adduced in support or criticism of a scientific hypothesis is always expressed in sentences, which frequently have the character of observation reports... the evidence . . . consists, in the last analysis, in data accessible to what is loosely called direct observation and such data are expressible in the form of "observation reports" '.'4 what are accessible to observation are objects and events. these are not the same things as facts nor, of course, as sentences, sentences can, in some loose sense, express facts (i.e. true propositions), and describe objects. none of what i say

<sup>10</sup> moritz schlick, the foundation of knowledge', *lp* 215.
 <sup>12</sup> ibid. 220.

<sup>13</sup> carl hempel, 'studies in the logic of confirmation', in *aspects of scientific explanation*, 14. <sup>14</sup> ibid. 21-2. is criticism of hempel's exemplary work on confirmation as a relation between sentences; i am using the distinction among various ways of describing positive instances of laws (or lawlike sentences) in order to emphasise the very different ways in which it is natural to talk of evidence.

perhaps it is not strange to call a black raven—some actual bird an *instance* of a law, but it does seem odd to say the bird is *evidence* for the law. at best this seems to be shorthand for saying it is the *fact* that this bird is a black raven that constitutes the evidence; or we could speak of the truth of the proposition, or of some appropriate sentence. so far, however, we have not touched on the epistemological issue, the question what it means for *someone* to have a reason to accept the law, to *possess* evidence. neither the existence of the black raven nor the truth of the proposition or sentence that says there is a black raven in itself gives anyone a reason to believe there is a black raven, much less a reason to believe all ravens are black. for someone to have a reason to believe all ravens are black, it is necessary for him to *believe*, for example, that here is a black raven.

we are off on a well-worn track. surely it is not enough simply to believe that here is a black raven; not enough, either, that the belief should also be true. for both of these conditions together do not add up to evidence unless the person has an adequate reason for holding the belief. if the reason must be another belief, we are faced by an infinite regress or a circle. a regress would make knowledge impossible, while a circle would lead to the difficulties of a pure coherence theory of knowledge. i'll come back to the latter in a moment.

at this point we come to the various attempts to find states of mind that bridge or eliminate the gap between sensation, where no question of truth can arise, and judgement, which is plausibly a source of evidence. quasi-sentences like 'black here now' have been proposed as expressing such states of mind. and perhaps we will be persuaded that there are such states of mind if we overlook the fact that the verb has been omitted (since putting it in would push things too far in the direction of judgement) and that words like 'here' and 'now' cannot be understood except as involving a reference to an agent. in any case, the attempt to base science on such states of mind is doomed, since no one has ever succeeded in showing how to base knowledge of an objective, common world on such 'evidence'. even schlick, who somehow hoped to back protocol sentences of the form 'a experiences black at time *t*' by whatever it is that is expressed by

'black here now' as said or thought by a at *t*, did not believe in a construction of science or a public world, with the construction based on 'immediate experiences'.

it should be obvious that no appeal to perception can clear up the question what constitutes a person's ultimate source of evidence. for if we take perception to consist in a sensation caused by an event in the world (or in the body of the perceiver), the fact of causality cannot be given apart from the sensation, and the sensation cannot serve as evidence unless it causes a belief. but how does one know that the belief was caused by a sensation? only further beliefs can help. if perception is expressed by locutions like 'a perceives that there is a black raven', then this can certainly serve as evidence.this does not solve the problem, it only transfers it to the concept of perception, since to perceive that there is a black raven is to be caused by a raven, and *in the right way*, to believe that there is a black raven.

one is struck, in reading early writings of the members of the vienna circle, by the embarrassed way in which they refer to experience, what is immediately given, what is directly observed. thus schlick says that all meaning goes back to ostensive situations, 'and this means, in an obvious sense, reference to "experience" or "possibility of verification" '.'<sup>5</sup> in 'the turning point in philosophy' he says the act of verification in which the path to the solution finally ends is always of the same sort: it is the occurrence of a definite fact that is confirmed by observation, by means of immediate experience';<sup>16</sup> but in 'positivism and realism' he expresses grave doubts about terms like 'the given' *(das gegebene),* and worries that if we use the word 'experience,' we will 'presuppose a distinction between what experiences and what is experienced'.<sup>17</sup>

there is, then, good reason to conclude that there is no clear meaning to the idea of comparing our beliefs with reality or confronting our hypotheses with observations. this is not, of course, to deny that there is an ordinary sense in which we perform experiments and note the results, or discover in our everyday pursuits that some of our beliefs are true and others false. what should be denied is that these mundane events are to be analyzed as involving

- <sup>15</sup> moritz schlick, 'meaning and verification', 148.
- <sup>16</sup> moritz schlick, 'the turning point in philosophy', lp 56.
- <sup>17</sup> moritz schlick, 'positivism and realism'. *lp* 84.

evidence that is not prepositional in character—evidence that is not some sort of belief. no wonder neurath and carnap were attracted to the idea of a coherence theory!

of course, coherence theories of knowledge are not without difficulties, and these, schlick was, as we have seen, quick to point out. let me pause here for a moment to make the obvious distinction between a coherence theory of knowledge and a coherence theory of truth. in his 1935 paper 'on the logical positivist's theory of truth' hempel had barely distinguished between the two; understandably, since he was not then aware of tarski's method for defining truth semantically. he therefore was at the time inclined to think the only sense we can make of the phrase 'sentence s is true' is 's is highly confirmed by accepted observation reports'.<sup>18</sup> but the concept of being highly confirmed by accepted observation reports belongs rather in the domain of epistemology; and when coupled with the idea that protocol statements 'may only be characterized by the historical fact' that they are accepted (i.e. believed true), leads directly to a coherence theory of knowledge. this is the theory to which schlick objected.

schlick agreed with neurath that protocols belong, in a general way, to the hypotheses of science. they are objective, and therefore intersubjectively understood and testable. they are *about* observations or experiences, but they don't attempt to *express* them. they take the form, roughly at least, of sentences like 'a saw a black raven at time t'. it is clear that one cannot be certain of the truth of such sentences—not even a at time t can be certain he is seeing a black raven; or, leaving the question of certainty aside, it is clear that anyone who judges such a sentence to be true may be wrong. further evidence is always relevant, and may come to outweigh the evidence of the moment. where schlick disagreed with neurath is on the question whether there are indisputable grounds on the basis of which we judge protocols to be true.<sup>19</sup>

the objection to neurath's coherence theory was the standard objection to all such theories: consistency is not enough, since it leaves no basis on which to choose between various and conflicting consistent theories. perhaps a theory of this kind banishes the last

<sup>&</sup>lt;sup>18</sup> carl hempel, 'studies in the logic of confirmation', in *aspects of scientific explanation*, 42.

<sup>&</sup>lt;sup>19</sup> moritz schlick. the foundation of knowledge'. lp 213. the following seven quotations are from lp 218-25.

remnant of absolutism from philosophy, as schlick said of neurath's claims, but it leaves us with no basis for judging truth. schlick insisted that we must have, and do have, indubitable grounds for choosing some sentences as the true ones rather than others. there are statements that are *not* protocol statements, which 'express facts of one's own "perception" ('or whatever you like to call it', he adds). schlick then admits that 'in spite of the fact that statements of this sort seem so simple and clear, philosophers have found themselves in a hopeless labyrinth the moment they actually attempted to use them as the foundation of all knowledge'. but he thinks we can steer clear of the familiar difficulties if we remember that 'one's *own* statements in the end play the only decisive role'.

one must admit that schlick's attempt to explain his view ends in obscurity. the observation 'sentences' that constitute the 'ultimate criterion' of all knowledge are not really sentences, being always of the form 'here now so and so'. such sentences cannot be written down (since they lose their certainty in a moment); they express a 'feeling of fulfillment, a quite characteristic satisfaction: we are *.satisfied'*. 'one cannot build any logically tenable structure upon such confirmations, for they are gone the moment one begins to construct.' finally, 'the occasion of understanding [observation statements] is at the same time that of verifying them: i grasp their meaning at the same time as i grasp their truth'. one can sympathize with neurath for rejecting this last step. but then one is left with a coherence theory.

hempel calls the neurath-carnap position a 'restrained' coherence theory. the reason is that neurath and carnap do provide us with a criterion for picking out one scientific theory from among the consistent ones. the criterion is that it is the consistent theory that maximizes agreement with the statements historically held true by 'mankind especially the scientists of our culture circle'.<sup>20</sup> in the end, protocol sentences have no pride of place: like any others they may be abandoned if they conflict with too much else we hold true.

thus it turns out, rather surprisingly, that both schlick and neurath held views that could be called 'restrained' coherence theories. they agreed that everything in the corpus of science, including protocol sentences, must be viewed as only tentatively established at any stage in the progress of science, and all sentences remain open

<sup>&</sup>lt;sup>20</sup> carl hempel, 'on the logical positivist's theory of truth'. 57.

to revision in the light of new evidence. they likewise agreed that when revision was called for, there were no strict rules for deciding where the revision should be made; it was a matter for 'decision'. their sole important difference concerned the question how the whole pattern of sentences accepted at a given time by science was to be related to experience, observation, or the real world. and on this score their answers were less than clear or satisfactory.

schlick's answer was unsatisfactory because it ended with something so private that even its meaning could only be given at a moment for an individual. how such a basis could warrant belief in a public objective world was not explained. neurath rejected the idea of a confrontation between a belief about the world and the world itself as well as the idea of an incorrigible subjective basis for scientific knowledge. but his suggestion as to how to 'restrain' a coherence theory is unappealing. he suggests that we start with protocols of the form 'a sees a black raven at t', and he dismisses the idea that such protocols are any more the basis of a's knowledge than of b's. this guarantees the intersubjective aspect of the language of science right down to the protocols (since, as neurath said, "every language ax such is inter-subjective'21-clearly a crack at schlick and his observation statements, whose meaning is revealed to only one person, and then only for a moment). neurath imagines all protocols being thrown into one great machine; a bell rings if a contradiction arises; something must then be thrown out, either one or more protocols, or perhaps a law or other theoretical statement; but 'who rebuilds the machine, or whose protocol sentences are thrown into the machine is of no consequence whatsoever'.22

there is an obvious difficulty here. if protocol sentences are known only by their form, throwing them *all* in the machine will mean each sentence *and* its negation will be thrown in, as well as 'a sees a black raven at spot *salt'* and 'b sees a non-black raven at spot *s* at f.' no basis for science can emerge from this, just endless consistent systems. if, on the other hand, the protocols are limited to the sentences that express *beliefs*, it *will* matter who mends the machine and whose protocols are thrown in or out. for each person will weigh the sentences he accepts (whether protocols or not) in accord with the strength of his beliefs—that's what it means to say they are *his* 

<sup>21</sup> otto neurath. 'protocol sentences', *lp*. 205. <sup>22</sup> [bid. 207.

beliefs. he will give weight to other people's protocols to the extent that he believes they are true, just as schlick maintained.

i would not dwell at such length on the familiar epistemological problems that beset schlick and neurath and their followers if i thought philosophy had now rejected them or solved them. on the contrary, i think the members of the vienna circle and their friends emphasized in a particularly useful way, even if partly in spite of themselves, a central unsettled problem in epistemology. and i think we can find a number of ideas and intuitions in the writings i have been discussing that point in the direction of a new view of the old problem.

the central problem may be stated as a dilemma in the theory of knowledge. each person has a complex network of beliefs. knowledge requires at least these two things: that some of these beliefs are true of the public world, and that each person has adequate reasons for holding these beliefs. i am willing to assume that the first condition is satisfied; the hard problem concerns the second issue, the way in which the system of beliefs is related to the world not merely semantically, but epistemologically. the dilemma arises because if we take as the connecting link something self-certifying (like schlick's observation statements or events), it is so private as to lack connection with the sentences of the public language that alone are capable of expressing scientific, or even objective, claims. but if we start with sentences or beliefs already belonging to the public language (or what can be expressed in it), we find no intelligible way to base it on something self-certifying (neurath's problem). in short, the foundations of knowledge must be subjective and objective at once, certain and yet open to question.

the problems i am rehearsing belong, we all know, to the foundations of epistemology, and in one form or another the problems are ancient. the logical positivists, one senses at once in their writings, were impatient with such problems, which they felt verged on the meaningless, or were to be solved by mere 'conventions' or 'decisions'. this attitude now seems to us, rereading these bold classics, cavalier. but the giddy conviction that a clear and correct line would somehow open up in the face of all that enthusiasm and intellectual power did produce some profoundly novel and valuable hints. what on the surface now may seem naive and failed attempts contained deeply suggestive intuitions of radical new ideas.

given the positivists' tendency to remain within the vocabulary of

intersubjective ideas and hypotheses, the flirtations of neurath, schlick, hempel, and carnap with some form of epistemological coherence theory is not surprising, but coherence theories have always been bedeviled by failure to distinguish between coherence theories of truth and coherence theories of knowledge. this is natural enough, since if knowledge, which is of the true, demands nothing but coherence of belief, how can truth require anything more than a set of coherent propositions? thus one is invited to dismiss the difference between coherence of beliefs and coherence of sentences (or statements or propositions) as tests of knowledge and truth, but while we find this confusion in the writings of schlick and neurath, we also find moments when the distinction is clearly made. neurath, we remember, hoped to destroy any aura of subjectivity in protocol statements by insisting that though it may be a 'historical accident' that one person is more inclined to accept his own protocols than those of someone else, in fact both are to be accepted on the same level, this leads him to the picture of the impersonal sorting machine into which protocol sentences are thrown. in proposing this idea, neurath never hints that it is only sentences held to be true by someone that are to count. it is this fact that occasioned schlick's outburst: 'the astounding error of the "coherence theory" can be explained only by the fact that its defenders and expositors were thinking only of such statements as actually occur in science.<sup>23</sup> here schlick means, of course, such as occur as assertions. yet in spite of much that we find in neurath's article on protocol sentences, can he really be accused of having forgotten the difference between the coherence of an arbitrary set of sentences and the coherence of a set of sentences held true? the famous metaphor of the ship which must be rebuilt at sea piece by piece proves he was aware that it is beliefs that are at stake, not mere sentences, for if we were dealing with an arbitrary set of sentences, nothing would stop us from putting them all in dry dock simultaneously, and hempel, as i mentioned above, makes the distinction clear. since standard objections to coherence theories of knowledge parallel standard objections to coherence theories of truth, it is not immediately apparent why it is so important to distinguish between them. but of course beliefs are not historically or causally arbitrary; even if our *reasons* for our beliefs are always other beliefs, the *causes* sometimes lie elsewhere. some appreciation of the

<sup>&</sup>lt;sup>23</sup> moritz schlick. the foundation of knowledge'. *lp* 215.
importance of this point must, as we will see, becredited to the logi cal positivists.

the logical positivists preferred to talk of sentences or 'statements' rather than beliefs, and we can easily enough make this switch as long as we remember that the sentences that correspond to beliefs are (1) sentences held true by someone, and (2) sentences that have an interpretation. someone else can know what i believe if he knows what sentences i hold true, and what those sentences mean. let me review some of the logical positivists' views in the light of these simple considerations.

there is schlick's idea that observation sentences are understood in the act of grasping their truth. this may well seem extreme or obscure, but it is related to the correct doctrine that an interpreter is constrained to take first person present tense attributions of attitude as presumptively true. such sentences (of english) as 'i believe i now see a black raven', if held true by a speaker, require of an interpreter that he assign a high a priori probability to their truth. this means: so interpret such sentences as to make them true when possible.

hempel comments on the fact that the protocol statements 'produced' by different men might not admit the construction of a unique system of scientific statements. he goes on: 'but fortunately this possibility is not realized: in fact, by far the greater part of scientists will sconer or later come to an agreement, and so, as an empirical fact, a perpetually increasing and expanding system of coherent statements and theories results from their protocol statements'.<sup>24</sup> again, we must assume that protocol statements are not any sentences written down, or uttered; they must be sentences their speakers believe to be true, or at least that a hearer believes the speaker to have held true. but it is surely odd to consider it merely 'fortunate' that there is a large degree of consensus: and why should we expect agreement to increase over time?

schlick has an even more surprising discussion of the possibility that someone might discover that all his own observations in no way substantiate the assertions made about the world by other men. he says that under these circumstances one would not, as neurath's protocol machine would, simply sacrifice one's own protocol statements. instead, one would cling to a

<sup>&</sup>lt;sup>24</sup> carl hempel. 'on the logical positivist's theory of truth', 57.

system of" knowledge into which one's own observations fitted unmutilated. and i can always construct such a system. i need only view the others as dreaming fools, in whose madness lies a remarkable method, or—to express it more objectively—i would say that the others live in a different world from mine . . . in any case no matter what world picture i construct, i would test its truth always in terms of" my own experience.<sup>25</sup>

this is a remarkable admission from someone who has objected that a coherence theory leaves us with an unacceptable relativism.

elsewhere in schlick, however, we find a rather different way of viewing the possibility of massive disagreement. he notes a basic contrast between a disagreement over whether two pieces of paper are the same color, and a disagreement over what color both pieces are. with respect to the first he says that 'by virtue of linguistic usage the proposition expresses just that experience'—i.e. the experience of sameness. but in the case of color, there is no objective way—that is, no way—to tell if you and i experience the same color agree entirely with mine, i cannot infer from this that you experience the same quality. as long as the inner order of your experiences agrees with mine, we will understand each other perfectly.<sup>26</sup>

schlick seems in the end to reject the view that the experiences may be undetectably different. '. . . the statement that different individuals have the same experience has its sole verifiable meaning in the fact that all their assertions . . . exhibit certain agreements . . . the statement *means* nothing but this.<sup>27</sup> it is not easy to tell from this passage whether schlick thinks the experiences might be qualitatively different while we could not in principle discern this, or meaningfully claim it; or whether he thinks no such situation could arise. the radical suggestion, which it is not impossible to read into schlick's attack on the coherence theory, is that interpersonal agreement, and hence objectivity, are built into the way in which we determine the meanings of other people's utterances, and hence the contents of their beliefs.

i mentioned above hempel's remark that 'fortunately' the protocol statements of different people allow the construction of a unique system of science. he adds that carnap has 'perhaps provided us a

" ibid. 93.

<sup>&</sup>lt;sup>25</sup> moritz schlick, the foundation of knowledge', *lp* 219.

<sup>&</sup>lt;sup>26</sup> here i paraphrase schlick, 'positivism and realism', *lp* 93.

possibility of explaining this fortunate fact'.<sup>28</sup> the possible explanation lies in the fact that 'young scientists are conditioned' to produce true protocol statements, and he adds, 'perhaps the fact of the general and rather congruous conditioning of scientists may explain to a certain degree the fact of a unique system of science.'<sup>29</sup>

it would be mysterious if people were first taught what various sentences mean, and *then* were conditioned to 'produce' the true ones; this would amount to teaching them, on the one hand, how to be better observers, and, on the other hand, how to be honest. but the situation may be seen rather as a matter of conditioning people, as we surely do, to hold certain sentences true under publicly observable conditions, and fixing on the interpretation of utterances of the sentences in accord with the success of the conditioning. this would explain interpersonal agreement on the main features of the environment in a natural way.

carnap at one point seems clearly to take this line. in 'psychology in physical language' (written in 1932) carnap flatly rejects neurath's idea that i must or can treat your protocol sentences on a par with my own.

generally speaking [he writes], a psychologist's spoken, written, or printed protocol sentences, when they are based on so-called introspection, are to be interpreted by the reader, and so figure in inter- subjective science, *not chiefly as scientific sentences, but as scientific facts.* the epistemological confusion of contemporary psychology stems, to a large extent, from this confusion of facts in the form of sentences with the sentences themselves considered as parts of science.<sup>50</sup>

the inferences we are permitted to draw from the fact that someone else utters a sentence are not the deductive consequences that flow from that sentence as interpreted, but rather the sort of inference we can draw from observing the movements of a voltmeter, or the movements of a raindrop. the point is not that others do not mean anything by the sentences they utter, but that we cannot take for granted that we know in advance *what* they mean; and interpretation is explicitly called for or implicitly assumed.

i think that by following out this line, along with several other suggestions drawn from passages i have been quoting, we can

<sup>&</sup>lt;sup>28</sup> carl hempel. 'on the logical positivist's theory of truth', 57.

<sup>&</sup>lt;sup>29</sup> ibid. 58.

<sup>&</sup>lt;sup>30</sup> rudolph carnap, 'psychology in physical language', *lp* 195.

discover the outline of a correct view of the foundations of empirical knowledge, a view that reconciles neurath's coherentist theory with schlick's insistence on a basic tie to experience and observation.

from here on, although i shall be borrowing in many and obvious ways on ideas of schlick, neurath, hempel, and carnap, i am stating my own position. this is a position deeply influenced by quine, though it is not quine's position.

neurath was right in rejecting the intelligibility of comparing sentences or beliefs with reality. we experiment and observe, but this is not 'comparing' in any but a metaphorical sense, for our experimentation bears no epistemological fruit except as it *causes* us to add to, cling to, or abandon our beliefs. this causal relation cannot be a relation of *confirmation* or *disconfirmation*, since the cause is not a proposition or belief, but just an event in the world or in our sensory apparatus. nor can such events be considered in themselves to be evidence, unless, of course, they cause us to believe something. and then it is the belief that is properly called the evidence, not the event.

neurath was also right in saying that, given this situation, we may as well admit that protocols, like any other propositions of science or common sense, can be wrong; we stand ready to tinker where tinkering does the most good. as hempel observed, no epistemological priority is left to protocols— they are like the rest. all this is, of course, the line quine was later to exploit in arguing against the analytic-synthetic distinction.

we are left, then, as neurath insisted, in a situation where our only evidence for a belief is other beliefs; this is not merely the *logical* situation, but also the pragmatic situation. and since no belief is self-certifying, no set of beliefs can supply a certain basis for the rest. how then can we escape schlick's objection that this makes 'arbitrary fairy stories to be as true as a historical report'? he concludes: 'thus the coherence theory is shown to be logically impossible . . . for by means of it i can arrive at any number of consistent systems of statements which are incompatible with one another.'<sup>31</sup> it's not clear what it means to say i could 'arrive' at various systems, since i do not invent my beliefs; most of them are not voluntary. still, the point of the criticism would seem to remain in the form of a challenge to say what reason i have to consider the bulk of my beliefs true.

<sup>&</sup>lt;sup>31</sup> moritz schlick. 'the foundation of knowledge', *lp* 216.

the key to the answer lies, i think, in generalizing carnap's two suggestions that we are conditioned to produce (hold true) specific sentences under particular conditions, and that we cannot use other people's statements as evidence until we have interpreted them. carnap said this only about protocol sentences, but the same should be said about all language.

language is in its nature, as neurath insisted, intersubjective; what someone else's words mean on a given occasion is always something that we can in principle learn from public clues. consider how we discover what some simple sentence means, say there's a table' or 'here's a piece of green paper'. our basic evidence is that the speaker is caused to assent (not just on this occasion, but generally) to these sentences by the presence of tables or pieces of green paper, while the absence of these objects causes him (generally) to dissent from the same sentences. i do not think of assent and dissent as overt speech acts, but as attitudes towards sentences sometimes revealed in speech and sometimes in other ways. my main point is that our basic methodology for interpreting the words of others necessarily makes it the case that most of the time the simplest sentences that speakers hold true are true. it is not the speaker who must perform the impossible feat of comparing his belief with reality; it is the *interpreter* who must take into account the causal interaction between world and speaker in order to find out what the speaker means, and hence what that speaker believes. each speaker can do no better than make his system of beliefs coherent, adjusting the system as rationally as he can as new beliefs are thrust on him. but there is no need to fear that these beliefs might be just a fairy tale. for the sentences that express the beliefs, and the beliefs themselves, are correctly understood to be about the public things and events that cause them, and so must be mainly veridical. each individual knows this, since he knows the nature of speech and belief. this does not, of course, tell him which of his beliefs and sentences are true, but it does assure him that his overall picture of the world around him is like the picture other people have, and is in its large features correct.

neurath, carnap, and hempel were right, i believe, in abandoning the search for a basic sort of evidence on which our knowledge of the world could rest. none is available, and none is needed. what they perhaps failed to appreciate is *why* it is not needed. it is not needed because the causal relations between our beliefs and speech, and the world also supply the interpretation of our language and of our beliefs. in this rather special sense, 'experience' is the source of all knowledge. but this is a sense that does not encourage us to find a mental or inferential bridge between external events and ordinary beliefs. the bridge is there all right—a causal bridge that involves the sense organs. the error lies, as neurath saw, in trying to turn this causal bridge into an epistemological one, with sense data, uninterpreted givens, or unwritable sentences constituting its impossible spans.

there are of course some beliefs that carry a very high degree of certitude, and in some cases their content creates a presumption in favor of their truth. these are beliefs about our own present propositional attitudes. but the relative certitude of these beliefs does not suit them to be the foundation of empirical knowledge. it springs, rather, from the nature of interpretation. as interpreters we have to treat self-ascriptions of belief, doubt, desire, and the like as privi leged; this is an essential step in interpreting the rest of what the person says and thinks. the foundations of interpretation are not the foundations of knowledge, though an appreciation of the nature of interpretation can lead to an appreciation of the essentially veridical nature of belief.

## 12 epistemology and truth

everyone agrees that what is known must be true, but beyond this there is little agreement on the epistemic status of truth. many philosophers in recent decades have held that truth is an epistemic concept; even when they have not explicitly held this thesis, their views have often implied it. coherence theories of truth are usually driven by an epistemic engine, as are pragmatic characterizations of truth, dummett and crispin wright's antirealism, peirce's idea that the truth is what science will end up believing, richard boyd's claim that truth is what explains the success of science, and hilary putnam's internal realism. quine also, at least at times, has maintained that truth is internal to a theory of the world and so to that extent dependent on our epistemological stance. relativism about truth is perhaps always a symptom of infection by the epistemological virus; this seems to be true in any case for quine, goodman, and putnam.

apparently opposed to all these views is the intuitive idea that truth, aside from a few special cases, is entirely independent of our beliefs; as it is sometimes put, our beliefs might be just as they are and yet reality—and so the truth about reality—be very different. according to this intuition, truth is 'radically non-epistemic' (so hilary putnam characterizes 'transcendental realism'), or 'evidencetranscendent' (michael dummett).

if we were to look for a short description of these two views of truth, we might well hit on the words 'subjective' and 'objective'; the assertion of an essential tie to epistemology introduces a dependence of truth on what can be verified by finite rational creatures, while the denial of such dependence makes truth objective. perhaps the words 'immanent' and 'transcendent' also suggest the same basic division.

and since the question of truth cannot be separated from questions about the nature of reality, 'antirealism' and 'realism' have often been used to make a closely related distinction. for reasons that will become clear as i go along, my preference is for 'subjective' and 'objective'; but i realize that some parties to the dispute will find this terminology tendentious.

in this talk 1 outline an approach to the concept of truth which i think undermines the dispute between the two views of truth. i do not aim to reconcile the two views, but to question their clarity, and so leave us free to give up the search for the 'right' theory of truth. that both views, while no doubt answering to powerful intuitions, are fundamentally mistaken is at least suggested by the fact that both invite skepticism. subjective theories are skeptical in the way idealism or many versions of empiricism are skeptical; they are skeptical not because they make reality unknowable, but because they reduce reality to so much less than we believe there is. objective theories, on the other hand, seem to throw in doubt not only our knowledge of what is 'evidence-transcendent', but all the rest of what we think we know, for such theories deny that there is any link between belief and truth.

let us consider first tarski-style truth definitions, with an eye to deciding whether they are objective or subjective or neither.<sup>1</sup> the first thing to be clear about is that tarski did not define, nor show how to define, the concept of truth; what he did was define truth for a specific language, and indicate how to define truth for other languages of a specific sort. tarski's basic insight was to make use of the apparently trivial fact that sentences of the form ' "snow is white" is true in 1 if and only if snow is white' must be true if the sentence quoted is a sentence of the language used to state the platitude. let us call such platitudes t-sentences. tarski observed that a definition of 'is true in I' which entailed a t-sentence for every sentence of I would pick out just the true sentences of I; and he further observed that the definition would be just as correct if the truth predicate were in a language that did not contain l, and the entailed t-sentences had a translation of the quoted sentence in place of that sentence itself to the right of the biconditional, the hard work came in providing a correct definition, given the constraints tarski imposed. (none of this survey is quite accurate, but my inaccuracies are irrelevant to the discussion.)

alfred tarski, the concept of truth in formali/ed languages'.

## epistemology and truth 179

it is clear that no definition in tarski's style is a definition of truth; each such definition is a definition of truth *in a particular language*. if we think that it is the same concept of truth that is involved in each of tarski's definitions, this can only be because we have a grasp of that concept that is independent of tarski's work.<sup>2</sup> in fact this is entirely clear; tarski's condition on satisfactory definitions of truth for particular languages rests on the recognition that t-sentences are obviously true, and this recognition depends on our prior understanding of the concept of truth. the triviality of t-sentences is a guarantee of their truth, and this guarantee is passed on by tarski's work to the 'partial definitions' of truth he provides. t-sentences are trivially true only if we know that the sentence described is identical with, or a translation of, the sentence that gives its truth conditions; but this is enough to display the sense in which we have a firm grasp of the general concept of truth.

tarski's definitions are sometimes called 'merely translational' or 'disquotational'. they are indeed disquotational in the special case where truth is being defined for a language in a language that contains it; in such cases, t-sentences show how to remove the quotation marks and get a sentence that provably has the same truth value. but this feature can be known to be present only when the object language is known to be contained in the metalanguage. the idea that such truth definitions are translational rests on a mistake. if one knows how to translate a language l into one's own language, it does not follow that one can automatically produce a truth definition. translation relates languages to one another; truth, even in tarski's limited editions, relates a language to the world.

tarski's definitions are reached through several steps. the notion of a well-formed expression (particularly a closed sentence) is defined; there is a recursive definition of a satisfaction relation (satisfaction is a highly generalized version of reference); truth is defined on the basis of the concepts of sentence and satisfaction; the recursive definitions are turned into explicit definitions. for present purposes it is best to forget the last step. what we are then left with is not a proper definition at all, but a recursive characterization of an undefined concept, that of satisfaction; and since truth is defined in terms of satisfaction, which is a semantic concept, it can no longer

<sup>&</sup>lt;sup>2</sup> as far as i know this point was first made by max black. the semantic definition of truth'. also see michael dummett. truth', in *truth and other enigmas*.

if my theory of truth, couched in my own words, is correct for my own language; this is, however, a point i shall not discuss here.

if we knew what makes a theory of truth true for a speaker or group of speakers, we could plausibly be said to understand the concept of truth; and if we could say what makes such a theory true, we could give an explicit account, if not a definition, of truth. the ultimate *evidence*, as opposed to a criterion, for the correctness of a theory of truth we must presume lies in available facts about how speakers use the language. when i say available, i mean publicly available—available not only in principle, but available in fact to anyone who is capable of understanding the speaker or speakers of the languages, all of us do understand some speakers of some languages, all of us must have adequate evidence for attributing truth conditions to the utterances of some speakers; all of us have, therefore, a competent grasp of the concept of truth as applied to the speech behavior of others.

have we now settled the question whether truth is radically nonepistemic, as objectivists claim, or basically epistemic, as subjectivists claim? it may seem so, since we have followed a course of argument that leads to the conclusion that it is how language is used that decides whether a theory of truth for that language is correct.

but in fact the matter is not settled, for it may be held by objectivists that the question whether the theory is true for a given language or group of speakers is indeed empirical, but only because the question of what the words mean is empirical; the issue of truth, it may be held, remains to be answered, whether by the theory itself or in some other way.

does the theory itself contain the answer? it does if a tarski-type theory of truth is a correspondence theory, for then the form of the theory in effect defines truth to be correspondence with reality—the classical form of objectivism with respect to truth. i am sorry to say that i have myself argued in the past that theories of the sort tarski showed how to produce were correspondence theories of a sort. i said this on the ground that there is no way to give such a theory without introducing a concept like reference or satisfaction which relates expressions to objects in the world.<sup>3</sup>

<sup>&</sup>lt;sup>1</sup> tarski seems to have viewed his truth definitions as implementing the idea of correspondence. i foolishly accepted the term in 'true to the facts', essay 3 in *inquiries into truth and interpretation*.

it now seems to me to have been a mistake to call such theories correspondence theories. let me explain why i think it was a mistake. the usual complaint about correspondence theories is that it makes no sense to suggest that it is somehow possible to compare one's words or beliefs with the world, since the attempt must always end up simply with the acquisition of more beliefs. this complaint was voiced, for example, by otto neurath, who for this reason adopted a coherence view of truth; carl hempel expressed the same objection, speaking of the 'fatal confrontation of statements and facts'.<sup>4</sup> richard rorty has repeatedly insisted, with dewey, that a correspondence view of truth makes the concept of truth useless.<sup>5</sup> i have said much the same myself.<sup>6</sup>

this complaint against correspondence theories i now think is a mistake. one reason it is a mistake is that it depends on taking the side of the subjectivists; it would be a legitimate complaint only if truth were an epistemic concept in the way subjectivists claim it is. if this were the only reason for rejecting correspondence theories, the objectivist can simply reply that his position is untouched; he always maintained that truth was independent of our beliefs or our abilities to learn the truth.

the real objection to correspondence theories is simpler; it is that there is nothing for true sentences to correspond to, the point was made long ago by c. i. lewis; he challenged the correspondence theorist to *locate* the fact or part of reality, or of the world, to which a true sentence corresponded. one can locate individual objects, if the sentence happens to name or describe them, but even such location makes sense relative only to a frame of reference, and so presumably the frame of reference must be included in whatever it is to which a true sentence corresponds. following out this line of thought led lewis to conclude that if true sentences correspond to anything at all, it must be the universe as a whole; thus all true sentences correspond to the same thing. frege, as we know, independently reached the same conclusion through a very similar course of reasoning. frege's argument, if alonzo church is right, can be formalized: starting from the assumptions that a true sentence can't be made to correspond to something different by the substitution of

<sup>&</sup>lt;sup>4</sup> see otto neurath. 'protocol sentences'; and carl hempel. 'on the logical positivist's theory of truth', 51.

<sup>&</sup>lt;sup>5</sup> richard rorty, *consequences of pragmatism*, introd.; also his 'pragmatism, davidson and truth'. <sup>6</sup> for example, in essay 10 in this volume.

coreferring singular terms, or by the substitution of logically equivalent sentences, one can show that if true sentences correspond to anything, they all correspond to the same thing. but this is to trivialize the concept of correspondence completely; there is no interest in the *relation* of correspondence if there is only one thing to correspond to since, as in any such case, the relation may as well be collapsed into a one-place predicate: 'x corresponds to the universe'. similarly, 'x corresponds to (or names) the true' or 'x corresponds to the facts' can less misleadingly be read 'x is true'. peter strawson made the point in his famous exchange on truth with j. l. austin. he observed that the parts of a sentence may correspond to parts of the world (that is, refer to them), but adds that there is nothing else in the world to which a sentence may be related. he correctly goes on to claim that 'while we certainly say that a statement corresponds to (fits, is borne out by, agrees with) the facts', this is merely 'a variant on saving it is true'.7

the correct objection to correspondence theories is not, then, that they make truth something to which we humans can never legitimately aspire; the real objection is rather that such theories fail to provide entities to which truth vehicles (whether we take these to be statements, sentences, or utterances) can be said to correspond. as i once put it, 'nothing, no *thing*, makes our statements true.' if this is right, and i am convinced it is, we ought also to question the popular assumption that sentences, or their spoken tokens, or sentencelike entities, or configurations in our brains can properly be called 'representations', since there is nothing for them to represent. if we give up facts as entities that make sentences true, we ought to give up representations at the same time, for the legitimacy of each depends on the legitimacy of the other.

i said i regretted ever having said that a tarski-style truth theory was a form of correspondence theory, and i do. my reason for saying it was not that i had made the mistake of supposing that sentences or utterances of sentences corresponded to anything. but i was still under the influence of the idea that truth is objective; the idea that truth, and therefore reality, are (except for special cases) independent of what anyone believes or can know. thus i advertised my view as a brand of realism, realism with respect to the 'external world', with respect to meaning, and with respect to truth.

<sup>7</sup> peter strawson,'truth', 194-5.

## epixtemology and truth 185

if the terms were ill chosen, it is because they suggest a positive endorsement of a position, or an assumption, that there is a clear positive thesis to be adopted. whereas all i was entitled to maintain, all that my position actually entailed with respect to realism and truth, was the negative claim that subjectivism is false. the objective view of truth, if it has any content, must be based on correspondence, correspondence as applied to sentences or beliefs or utterances, entities that are propositional in character, and such correspondence cannot be made intelligible. to the extent that realism is just the ontological version of a correspondence theory, i must also reject it. all i had in mind in calling my position a form of realism was a rejection of antirealism; i was concerned to reject the doctrine that either reality or truth depends directly on our epistemic powers. once more, however, i must make clear that i am neither accepting nor rejecting the objectivist-realist slogan that the real and the true are 'independent of our beliefs'. the only evident positive sense we can make of this phrase that consorts with the intentions of those who prize it derives from the idea of correspondence, and this idea must be rejected. but to reject it is not to say there is no connection whatever between belief and truth; there must be some connection, as we have seen, if we are to relate the truth of utterances to their human context. the question is what that connection may be.

various forms of subjectivism—that is, of views that make truth out to be an epistemic concept—connect human thoughts, desires, and intentions to truth in quite different ways, and i cannot pretend to do justice to all such views. the best i can do is explain why, despite the differences, it makes sense to group many of them together, and to indicate the reasons for being dissatisfied with them.

i have classified coherence theories as subjective, and this needs an explanation. a pure coherence theory of truth would hold, i suppose, that all the sentences in any consistent set of sentences are true. i doubt if anyone ever held such a theory, for it is obviously false. those who have proposed coherence theories, for example neurath and carnap (at one time), have usually made clear that it was sets of *beliefs*, or of sentences held to be true, whose consistency was enough to make them true; this is why i class coherence theories with other subjective views; they tie truth directly to what is believed. but unless something further is added, this view seems as

wrong as schlick held it to be (he called it an 'astounding error'<sup>8</sup>). the obvious objection is that many different consistent sets of beliefs are possible which are not consistent with one another.

there are theories not usually, or perhaps properly, thought of as coherence theories which have much the same drawback, quine holds that the truth of some sentences, which he calls observation sentences, is tied directly to sensory input; further sentences derive their empirical content from their connections with observation sentences and their logical relations to one another, the truth of these further sentences (call them theoretical sentences) depends on nothing but how well they serve to explain or predict observation sentences on the basis of other observation sentences, quine holds (plausibly) that there could be two theoretical structures which are equally capable of accounting for all true observation sentences, and yet such that neither can be reduced to the other (each theory contains at least one predicate that cannot be defined using the resources of the other theory), quine has at different times embraced two different ways of thinking of this situation. according to one thesis, both theories are true. to this i have no objection. according to the other view, a speaker or thinker at a given time operates with one theory, and for him at that time the theory he is using is true and the other theory false. if he shifts to the other theory, then *it* is true and the other theory false. this position seems to be what quine has meant on occasion when he has said that truth is 'immanent'.<sup>y</sup> this conception of the immanence or relativity of truth should not be confused with the pedestrian sense in which the truth of sentences is relative to the language in which they occur. quine's two theories can belong to, and be stated in, the same language. it is not easy to see how the same sentence (without indexical elements), with meaning unchanged, can be true for one person and not for another, or for a given person at one time and not at another, the difficulty seems due to the attempt to import epistemological considerations into the concept of truth.

putnam's 'internal realism' also makes truth immanent, though not, as quine's view does, to a theory, but to the entire language and

<sup>&</sup>lt;sup>8</sup> moritz schlick, "the foundations of knowledge'.

<sup>&</sup>lt;sup>9</sup> w. v. quine, *ontological relativity and other essays*. for quine's problem about empirically equivalent, mutually irreducible theories, see his 'on empirically equivalent systems of the world', 313-28, *theories and things*, 29-30. and quine in hahn and schilpp (eds.). *the philosophy of w. v. quine*, 156-7.

## epistemology and truth 187

conceptual scheme a person accepts. of course, once again, if all that this means is that the truth of sentences or utterances is relative to a language, that is familiar and trivially correct. but putnam seems to have something more in mind, for example that a sentence of yours and a sentence of mine may contradict each other, and yet each be true 'for the speaker'. it is hard to think in what language this point can be persuasively stated. the source of the trouble is once again the felt need to make truth accessible. putnam is clear about this. he identifies truth explicitly with idealized justified assertability. he calls this a form of realism because there is 'a fact of the matter as to what the verdict would be if the conditions were sufficiently good, a verdict to which opinion would "converge" if we were reasonable'. he adds that his view is 'a human kind of realism, a belief that there is a fact of the matter as to what is rightly assertable for us, as opposed to what is rightly assertable from the god's eye view so dear to the classical metaphysical realist'.<sup>10</sup> one suspects that if the conditions under which someone is ideally justified in asserting something were spelled out, it would be apparent either that they allow the possibility of error or that they are so ideal as to make no use of the intended connection with human abilities. it is also striking that putnam seems to have no argument for his position except that the alternative ('metaphysical realism', i.e. a correspondence theory) is unacceptable. he does not argue that there can be no other position.

putnam describes his position as close to michael dummett's on the main point, the epistemological status of truth. one difference is that putnam is less certain than dummett that truth is limited to what is definitely ascertainable, and therefore he is less sure that the principle of bivalence must be abandoned; this perhaps explains why putnam calls his view a form of realism while dummett calls his position antirealism. putnam also thinks he differs from dummett in tying truth to *idealized* justified assertability instead of justified assertability; but here i think a close reading of dummett would show that he has much the same idea; certainly crispin wright does. if dummett does not insist on something similar to putnam's ideal conditions, then i think a criticism of dummett that putnam once formulated applies: if truth depends simply on justified assertability, truth can be 'lost', and this must be wrong." dummett says he agrees

<sup>&</sup>lt;sup>10</sup> hilary putnam, realism and reason, p. xviii.

<sup>&</sup>lt;sup>11</sup> hilary putnam. 'reference and understanding' and 'reply to dummett's comment', in margarit (ed.). *meaning and use.* 

that truth can't be lost, but he fails to give a clear idea of how warranted assertability can be both a fixed property *and* one that depends on the actual ability of human speakers to recognize that certain conditions are satisfied. actual abilities wax and wane, and differ from person to person; truth does not.

why does durnmett endorse this view of truth? there are a number of reasons, but one seems to be this: we have seen that a theory of truth in tarski's style neither defines nor fully characterizes truth; there is no way to tell if the theory applies to a speaker or group of speakers unless something is added to relate the theory to the human uses of language. dummett thinks the way (perhaps the only way) to do this is to make truth something humanly recognizable. the human use of language must be a function of how people understand the language, so if truth is to play a role in explaining what it is to understand a language, there must be something that counts as a person having 'conclusive evidence' that a statement is true. one can appreciate the force of this idea while finding it difficult to accept. i have given my chief reason for rejecting it; that it is empty, or makes truth a property that can be lost, but there are other strong intuitions that would also have to be sacrificed if dummett were right. one is the connection of truth with meaning: on dummett's view we can understand a sentence like 'a city will never be built on this spot' without having any idea what it would be for this sentence to be true (since the sentence, or an utterance of it, has no truth value for dummett). another is the connection of truth with belief: on dummett's view i can understand and believe a city will never be built on this spot, but my belief will have no truth value.

i might be tempted to go along with dummett if i thought we must choose between what putnam calls transcendental realism, i.e. the view that truth is 'radically non-epistemic', that all our best researched and established beliefs and theories may be false, and dummett's identification of truth with warranted assertability, since i find the former view, essentially the correspondence view, incomprehensible, while i find dummett's view merely false.

but i see no reason to suppose that realism and antirealism, explained in terms of the radically nonepistemic or the radically epistemic character of truth, are the only ways to give substance to a theory of truth or meaning. let me suggest a way that seems to me different from either of the two ways we have been discussing. first, i assume that there are inescapable and obvious ties among the

## epistemology and truth 189

concepts of truth, belief, and meaning. if a sentence s of mine means that p, and i believe that p, then i believe that s is true. what gives my belief its content, and my sentence its meaning, is my knowledge of what is required for the belief or the sentence to be true. since belief and truth are related in this way, belief can serve as the human attitude that connects a theory of truth to human concerns.

we observed that the evidence for the correctness of a theory of truth must come at the level of the t-sentences; these also constitute the testable predictions of the theory, the evidence that the utterances of a certain sentence by a speaker are true under the conditions provided by a t-sentence consists in facts about the conditions that cause the speaker to hold the sentence true and to take that sentence to express a belief of his. of course speakers sometimes hold sentences true when they are not; these sentences express mistaken beliefs. so the relation between belief and truth is not simple. but that it exists, and is central to meaning and truth, can be seen from the simplest cases, and their role in understanding other cases. the simplest cases are those where a sentence such as that's a book' or this is yellow' is caused to be held true by the conspicuous presence of books or yellow things. this is evidence that these sentences are true just when books or yellow things are present; the reason is that what determines the meaning of such sentences is what routinely makes them true. clearly this account must be elaborated in many ways to account for the possibility of error; much more must be said about sentences more remote from direct observation; provision must be made for the truth conditions of sentences held true no matter what is observed; and so forth. i have explored these matters elsewhere. the important matter in the present context is that it is the role of belief that connects truth and meaning; and if this is so, truth is neither radically nonepistemic, nor radically epistemic.

the approach to truth that i recommend is not epistemic in putnam's or dummett's sense, nor in the sense of peirce, or james, or dewey, for it makes no attempt to tie truth in general to what can be ascertained by rational creatures. but neither is this approach radically nonepistemic, for it does not follow from my view that all our best-attested and widely held beliefs may be false. by paying close attention to what gives content to our beliefs and meaning to our sentences, we see that it is impossible for most of our perceptual beliefs to be false. what we hold true, what we believe, determines what we mean, and thus, indirectly, when our sentences are

true. believing doesn't make it so, but it creates a presumption that it is so.

i have two further observations, and a conclusion. the first observation is that holding a sentence true, or taking it to express a belief, is not a *use* of language, and so it may seem that it cannot provide evidence for what people mean by what they say. the conclusion does not follow. there are obvious relations between holding a sentence true and linguistic (and other) behavior. if we could not often fathom from his linguistic behavior when a speaker held his sentences true, we could not interpret his speech. i have chosen the attitude of belief rather than acts of assertion as providing the central clue to truth not because i deem the actual use of language unimportant, but because i know of no simple relation between the intentions with which we utter or write sentences, and what our words mean. our speech acts reveal our underlying attitudes towards our sentences; but often indirectly.

the second observation is that my approach to the concept of truth, like my approach to questions about meaning, assumes that what we want is an account of what it is to understand a language. in this i agree with putnam, quine, dummett, and others. but in one important respect quine and i differ from dummett, and perhaps putnam (though i am less sure about putnam). dummett, at any rate, thinks of an account of what it is to understand a language as demanding a description of what it is to grasp the meanings of expressions. i have never been sure i understood this metaphor, and have therefore asked instead what it would suffice someone to know if he were to understand an arbitrary utterance of a speaker (my answer: a theory of truth for that speaker's language), plus an account of how someone could come to know that such a theory was true. dummett's approach looks at language from the subjective point of view, the perspective of the speaker. my approach looks at language from the start as a social transaction and therefore concentrates on what one person can learn about another in the context of a shared world.

the conclusion i promised is this. all attempts to characterize truth that go beyond giving empirical content to a structure of the sort tarski taught us how to describe are empty, false, or confused. we should not say that truth is correspondence, coherence, warranted assertability, ideally justified assertability, what is accepted in the conversation of the right people, what science will end up maintaining, what explains the success of science or of our ordinary beliefs.

## epistemology and truth 191

to the extent that realism and antirealism depend on one or another of these views of truth we should refuse to endorse either. realism, with its insistence on radically nonepistemic correspondence, asks more of truth than we can understand; antirealism, with its limitation of truth to what can be ascertained, unnecessarily deprives truth of its role as an intersubjective standard. if we want to speak the truth about truth, we should say no more than need be.

# 13 epistemology externalized

we know in a way no one else can what we believe, fear, want, value, and intend. we know how things seem to us, how they look to us, feel to us, smell and sound to us to be. we know these things in a way we can never know about the world around us. whether or not we are sometimes wrong about the contents of our own minds, whether or not we can be in doubt about our own sensations and thoughts, one thing is certainly true of such beliefs: they cannot be generally mistaken. if we think we have a certain thought or sensation, there is a strong presumption that we are right.

the special authority that pertains to first person knowledge has seemed to many philosophers to suit it uniquely to the task of supplying a foundation for the rest of knowledge, particularly, of course, for our knowledge of the 'external world' and of the minds of others. such knowledge stands in need of a foundation, it is thought, precisely because there is no presumption that our beliefs about the world or the minds of others are true.

i have no intention of rehearsing the difficulties this picture of the mind and its place in nature has gotten us into, nor of reminding you of the ingenious but unpersuasive ways that have been proposed for circumventing its obvious failings. in this essay i describe an alternative, an alternative that i think is correct.

quine has proposed that epistemology be naturalized.<sup>1</sup> by this he has meant that philosophy should abandon the attempt to provide a foundation for knowledge, or otherwise to justify it, and should instead give an account of how knowledge is acquired. critics have complained that in abjuring the traditional normative task of episte-

<sup>&</sup>lt;sup>1</sup> w. v. quine, 'epistemology naturalized', in *ontological relativity and other essays*.

mology quine has simply changed the subject; i suspect that he agrees, and that this is part of what he intended. of course, the distinction between describing and justifying, between an empirical account of the genesis of knowledge and a statement of the norms belief must satisfy to count as knowledge, is by no means clear, as is perhaps evident from the fact that quine's naturalized epistemology, while it makes no serious attempt to answer the skeptic, is recognizably a fairly conventional form of empiricism.<sup>2</sup> and in any case, how can one describe knowledge or its origins without deciding what knowledge is? quine's answer to this question is that we must accept what science and enlightened common sense dictate without trying to justify it; his account of how we came by this knowledge is, however, just the kind of account that has traditionally been taken to constitute an attempt at justification.

i do not accept quine's account of the nature of knowledge, which is essentially first person and cartesian. but i do find congenial quine's resolutely third person approach to epistemology, and to the extent that the naturalization of epistemology encourages or embraces such an approach, i am happy to count myself a naturalized epistemologist.

there is at least a presumption that we are right about the contents of our own minds; so in the cases where we are right, we have knowledge. but any particular item of such knowledge is logically independent of our beliefs about a world outside, and so cannot supply a foundation for science and commonsense beliefs. this is how skeptics, like hume, reason, and i think they are right: knowledge of the contents of our own minds cannot be the basis for the rest of our knowledge. if this is correct, then our beliefs about the world must, if they are to count as knowledge, stand alone. yet it has seemed obvious to many philosophers that if each of our beliefs about the world, taken alone, may be false, there is no reason why all such beliefs might not be false.

this reasoning is fallacious. it does not follow, from the fact that any one of the bills in my pocket may have the highest serial number, that all the bills in my pocket may have the highest serial number, or from the fact that anyone may be elected president, that everyone may be elected president. nor could it happen that all our beliefs

<sup>2</sup> for this characterization of quine's epistemology, see essay 10 in this volume, and my 'meaning, truth and evidence'.

## epistemology externalized 195

about the world might be false. suppose i think i see a mouse disappear behind a chair. clearly this belief could be mistaken. but would this belief" be wrong if i did not truly believe a mouse was a small, four-footed mammal, or a chair an object made for sitting? perhaps. there may be no saying exactly what other true beliefs i must have in order to have a particular false belief. but it seems clear that a belief of any kind, true or false, relies for its identification on a background of true beliefs; for a concept, like that of mouse or chair, cannot remain the same concept no matter what beliefs it features in. it is possible to try to avoid this conclusion by arguing that a belief such as that a mouse is a small, four-footed animal is true by virtue of the concepts alone-it is an analytic truth-and so is not really about the world. one could then still say all our beliefs about the world might be false. this line is not available to someone who, like me, does not think a clear line between analytic and synthetic truths is there to be drawn, but even if there are indubitably analytic truths, it is not plausible that these serve to eliminate the host of cases in which concepts are individuated by multiple empirical criteria.

because of the holistic character of empirical belief, then, it is impossible that all our beliefs about the world are false. reflection on the reasoning that leads to this thesis suggests, however, that it has disturbing limits. what it shows is that we cannot harbor particular false (or true) beliefs about individual objects unless we have many true beliefs about the nature of such objects. this leaves open the possibility that we may be wrong in all our particular beliefs about what exists in the world, and this would be a pretty extreme skepticism, though not quite total.

but could this happen? suppose i know what a mouse is and what a chair is. how did i arrive at my false belief that i saw a mouse disappear behind a chair? presumably i interpreted what i saw, or experienced, as i did because it satisfied my criteria for the presence of a mouse and of a chair (and of disappearing). but we have just decided that i cannot be generally mistaken about the criteria. so my mistake could have come about only in this way: the scene that presented itself to me was one that justified my belief, even though the belief was false. this can certainly happen. it is hard, though, to see how it could happen always, or even often. for how could the scenes that present themselves to me justify beliefs none of which are true?

no skeptic will be persuaded by this argument. i will try to make

it more persuasive. allow me for the moment to make an assumption that begs the question of an external world inhabited by other people; let us suppose we have a speaker and someone else who is trying to understand his words. each time a mouse appears nearby in good light and with the speaker oriented in the direction of the mouse, etc., the speaker utters what sounds to the interpreter like the same expression: 'raton'. when the lighting is poor or the speaker inattentive, the response is less firmly correlated with mouse appearances. i think that unless there is a host of evidence against such an interpretation, the competent interpreter will take the speaker to mean by his words, and to believe, that there is a mouse present. what recommends this interpretation is the fact that the presence of a mouse has apparently in each case caused the speaker intentionally to utter the same expression, 'raton', and to utter it in an affirmative spirit. of course an interpreter may make a mistake: he may be in error about the intention of the speaker, or about the true regular and exclusive cause of the speaker's response, the interpreter would be wrong, for example, if a squirrel would provoke the same response in the speaker, for then mice would not be the regular and exclusive cause of utterances of 'raton'. but provided the interpreter makes none of these mistakes (nor certain others to be mentioned), his interpretation is correct.

suppose the speaker sees a skillfully made mechanical mouse and sounds his 'raton'. the interpreter must decide whether the speaker's word and concept include mechanical mice, or whether this is a mistake on the speaker's part. deciding this is not beyond the interpreter's power. he may expose the mechanism to the speaker, for example, and note the response. the careful interpreter will have to learn how the speaker responds to many other classes of objects before he arrives at his final interpretive scheme.

obviously the matter is subtle and complicated. but as long as we adhere to the basic intuition that in the simplest cases words and thoughts refer to what causes them, it is clear that it cannot happen that most of our plainest beliefs about what exists in the world are false. the reason is that we do not first form concepts and then discover what they apply to; rather, in the basic cases, the application determines the content of the concept. an interpreter who starts from scratch—who does not already understand the language of a speaker—cannot independently discover what an agent's beliefs are about, and then ask whether they are true. this is because the situations which normally cause a belief determine the conditions in which it is true.

these remarks apply directly to some beliefs only. someone who has learned from books what a guanaco looks like may never have been caused to accede to 'that's a guanaco' by seeing a guanaco, and yet be prepared (having seen pictures of guanacos perhaps) to accede when he does see one. or, to take a harder case, someone may know, in some reasonable sense, what a guanaco is, and that it is not a llama, and yet be regularly caused to assent to 'that's a guanaco' in the presence of llamas. in both these cases, the contents of the belief that there is a guanaco present is determined, not by exposure to guanacos, but by having acquired other words and concepts, such as those of llama, animal, camel, domesticated, and so forth. somewhere along the line, though, we must come to the direct exposures that anchor thought and language to the world.

the general approach to the identification of mental states that i am urging has something in common with what is now sometimes called externalism or (by tyler burge) anti-individualism; but it also differs from the best-known versions of externalism in important respects. so at this point i want to distinguish my version of externalism from other forms of externalism, particularly those of hilary putnam and tyler burge.

i return to first person authority, the fact that each person generally knows what he thinks without appeal or recourse to evidence, and thus knows what is in his own mind in a way that no one can know what is in the mind of another person. it has been widely supposed that externalism, which holds that the contents of a person's propositional attitudes are partly determined by factors of which the person may be ignorant, cannot be reconciled with first person authority. thus putnam decided that because (a form of) externalism is true, meanings 'just ain't in the head'—and held that the same applied to thoughts generally.<sup>3</sup> others have come to the same conclusion.<sup>4</sup> the conclusion does not follow, at least for the kind of externalism i have described. i have argued this at length elsewhere,<sup>5</sup> but the reasoning can be summarized as follows. an interpreter must discover, or correctly assume on the basis of indirect

in essay 2 in this volume.

<sup>&</sup>lt;sup>1</sup> hilary putnam, 'the meaning of "meaning" ', 227.

<sup>&</sup>lt;sup>4</sup> for examples, see john searle. *intentionality*; and andrew woodfield. 'introduction', in woodfield (ed.), in *thought and object*, p. viii.

evidence, what the external factors are that determine the content of another's thought; but since these factors determine both the contents of one's thought and the contents of the thought one believes one has (these being one and the same thought), there is no room for error about the contents of one's own thought of the sort that can arise with respect to the thoughts of others.

while i do not think that the main form of externalism that putnam endorses<sup>6</sup> threatens first person authority, i am not entirely happy with his thesis for other reasons. putnam's externalism applies mainly to natural kind words like 'water' and 'leopard'. the idea is that if i learn the word 'water' while experiencing h2o, the word must henceforth refer only to substances with the same microstructure. while i agree, as i said above, that the usual cause of my use of the word determines what it means, i do not see why sameness of microstructure is necessarily the relevant similarity that determines the reference of my word 'water'. (i will say what i think does control the relevant similarity presently.) i also think there is no reason to limit externalism to one, or a few, categories of words. i believe it is characteristic of language and thought quite generally that their ties to the world accrue from the sort of causal connections i have been discussing.

tyler burge has argued for two forms of externalism. in his earlier articles<sup>7</sup> he concentrated on the idea that the meanings of a person's words, and the contents of that person's thoughts, depend in part on the linguistic practices of the person's community, even in cases where the individual is mistaken about the relevant practices. later articles<sup>8</sup> emphasized the ways in which the contents of utterances and thoughts depend on the causal history of the individual, particularly in connection with perception. let me call the two brands of externalism social externalism and perceptual externalism. as burge has argued for them, i think they are largely independent in the sense that neither entails the other.

i am not impressed with the arguments for social externalism, for three reasons.

*first,* it seems to me false that our intuitions speak strongly in

<sup>&</sup>lt;sup>6</sup> i have in mind the sort of externalism inspired by twin earth examples rather than those prompted by the 'linguistic division of labor'.

<sup>&</sup>lt;sup>7</sup> tyler burge, 'individualism and the mental', 'other bodies', two thought experiments reviewed', and 'individualism and psychology'.

<sup>&</sup>lt;sup>8</sup> e.g. 'cartesian error and the objectivity of perception', 'intellectual norms and foundations of mind', and "individualism and self-knowledge'.

favor of understanding and interpreting an agent's speech and thoughts in terms of what others would mean by the same words. for one thing, there is the problem of deciding what group is to determine the norms. but more important, we understand a speaker best when we interpret him as he intended to be interpreted; this will explain his actions far better than if we suppose he means and thinks what someone else might mean and think who used the same words 'correctly'.

*second, i* think there is a conflict between surge's social externalism, which ties a speaker's meaning to an elite usage he may not be aware of, and first person authority.

*third*, i have a general distrust of thought experiments that pretend to reveal what we would say under conditions that in fact never arise. my version of externalism depends on what i think to be our actual practice.<sup>9</sup>

i turn now to burge's perceptual externalism. he makes two important points. the first is that we can have perceptual knowledge, say that we are seeing a cow, without having to ascertain independently that we are not deceived in one way or another (a horse painted to look like a cow, a hologram, etc.). when we have such a case of perceptual knowledge, the content of our thought is partly determined by the cause of the thought. here is what burge says:

to think of something as water . . . one must be in some causal relation to water—or at least in some causal relation to other particular substances that enable one to theorize accurately about water. in the normal cases, one sees and touches water. such relations illustrate the sort of conditions that make possible thinking of something as water . . . to *think* that water is a liquid, one need not *know* the complex conditions that must obtain if one is to think that thought.<sup>10</sup>

it is hard to decide, from this passage, whether the contents of the thought that this is water are determined, in the case where the thought is true, by the fact that it is water that is causing the thought, or whether in this case, as in the cases where the thought is false, the contents are determined by the normal case, the general run of the thinker's causal contact with his environment. i see no obvious difficulty in accepting and combining both these views, though of course there must be a distinction between the role the actual presence of

<sup>&</sup>lt;sup>9</sup> see essay 2 in this volume. <sup>10</sup> tyler burge, 'individualism and self-knowledge', 653-4.

water plays when it causes me to think this is water, and the role the history of my relations to water play in making a false thought that i am seeing water nevertheless a thought about water. clearly it is the latter thesis that is essential to externalism. burge puts it this way: the natures of such states are determined partly by normal relations between the person . . . and the environment. error is determined against a background of normal interaction.'11

i am in sympathy both with the epistemological view that perceptual knowledge does not require that we know independently that the enabling conditions obtain, and of course with the view (more important in the present context) that the contents of our thoughts and sayings are partly determined by the history of causal interactions with the environment.

agreeing on this point comes naturally to someone like me who has for some thirty years been insisting that the contents of our earliest learned and most basic sentences ('mama', 'doggie', 'red', 'fire', 'gavagai') must be determined by what it is in the world that causes us to hold them true. it is here, i have long claimed, that the ties between language and the world are established and that central constraints on meaning are fixed; and given the close connections between thought and language, analogous remarks go for the contents of the attitudes.

burge and i are essentially in agreement, then, on perceptual externalism, but there is a closely related matter on which we differ, and this concerns the question with which i began this essay, the question whether externalism provides an answer to certain sorts of skepticism. burge says, 'there is no easy argument against skepticism from anti-individualism and authoritative self-knowl-edge'.<sup>12</sup> in one sense, i agree, but only because i don't see what first person authority has to do with it. but it does seem to me that *if* you accept perceptual externalism, there *is* an easy argument against global skepticism of the senses of the sort that descartes, hume, russell, and endless others have thought requires an answer.

most perceptual representations are formed and obtain their content through regular interaction with the environment. they represent what, in some complex sense of 'normally', they normally stem from and are applied to. it

<sup>&</sup>lt;sup>11</sup> 'cartesian error and the objectivity of perception'. 125.

<sup>&</sup>lt;sup>12</sup> 'individualism and self-knowledge'. 655 n.

makes no sense to attribute systematic perceptual error to a being whose perceptual representations can be explained as the results of regular interaction with a physical environment  $\dots^n$ 

it may seem that this assumes more than mere externalism. but i don't see how. burge considers a case much like standard brain-inthe-vat cases and shows why such a brain cannot (for long) be radically deceived about its environment. i approve of this argument, having often used it myself. if anything is systematically causing certain experiences (or verbal responses), that is what the thoughts and utterances are about. this rules out systematic error. if nothing is systematically causing the experiences, there is no content to be mistaken about. to quote myself: 'what stands in the way of global skepticism of the senses is, in my view, the fact that we must, in the plainest and methodologically most basic cases, take the objects of a belief to be the causes of that belief.<sup>14</sup> anyone who accepts perceptual externalism knows he cannot be systematically deceived about whether there are such things as cows, people, water, stars, and chewing gum. knowing why this is the case, he must recognize situations in which he is justified in believing he is seeing water or a cow. in those cases where he is right, he knows he is seeing water or a cow.

i have rejected social externalism as championed by burge. do i therefore think social factors play no role in the externalism of the mental? not at all. but i would introduce the social factor in a way that connects it directly with perceptual externalism, thus locating the role of society within the causal nexus that includes the interplay between persons and the rest of nature.

burge and i agree that the cause of certain mental states is relevant to the content of those states. and we agree that one kind of case is especially important: an example is the way the fact that a certain mental state has been typically caused by seeing cows allows us to think 'there's a cow' even when no cow is present. but here a problem arises. what determines the content of such basic thoughts (and what we mean by the words we use to express them) is what has typically caused similar thoughts. but what *has* typically caused them? there are many choices, for example events that occurred before all cows, or events spatially closer to the thinker than any cow. burge at

<sup>14</sup> essay 10 in this volume.

<sup>&</sup>lt;sup>13</sup> 'cartesian error and the objectivity of perception', 131.

one point says that causes antecedent to the 'natural' cause or intermediate between the natural cause and the mental effect ('such as arrays of light striking the retina') would make the description of the content 'complicated in ways that have never been fully articulated'.15 but complicated for whom? articulated by whom? of course all those arrays of light striking the retina have something in common; and so for other classes of causes. it is we humans for whom these classifications are complicated and impossible to articulate (except, of course, by cheating in familiar ways). it is we who class cow appearances together, more or less naturally, or with minimal learning, and even so, another classification is required to complete the point, for the class of relevant causes is in turn defined by similarity of responses: we group together the causes of someone's responses, verbal and otherwise, because we find the responses similar. what makes these the relevant similarities? the answer again is obvious; it is we, because of the way we are constructed (evolution had something to do with this), who find these responses natural and easy to class together. if we did not, we would have no reason to claim that others were responding to the same objects and events (i.e. causes) that we are. it may be that not even plants could survive in our world if they did not to some extent react in ways we find similar to events and objects that we find similar, this clearly is true of animals; and of course it becomes more obvious the more like us the animal is.

the identification of the objects of thought rests, then, on a social basis. without one creature to observe another, the triangulation that locates the relevant objects in a public space could not take place. i do not mean by this that one creature observing another provides either creature with the concept of objectivity, the presence of two or more creatures interacting with each other and with a common environment is at best a necessary condition for such a concept. only communication can provide the concept, for to have the concept of objectivity, the concepts of objects and events that occupy a shared world, of objects and events whose properties and existence is independent of our thought, requires that we are aware of the fact that we share thoughts and a world with others.

for this reason we cannot resolve the question of the contents of

<sup>&</sup>lt;sup>15</sup> 'cartesian error and the objectivity of perception', 126-7.

## epistemology externalized 203

mental states from the point of view of a single creature. this is perhaps best seen by thinking about how one person learns from another to speak and think of ordinary things. put in greatly simplified terms, a basic aspect of such learning can be described in this way: the learner is rewarded, whether deliberately or not, when the learner makes sounds or otherwise responds in ways the teacher finds appropriate in situations the teacher classes together. the learner is subsequently caused to make similar sounds by situations the learner instinctively classes together, corrections are possible, of course. success at the first level is achieved to the extent that the learner responds with sounds the teacher finds similar to situations the teacher finds similar. the teacher is responding to two things: the external situation and the responses of the learner. the learner is responding to two things: the external situation and the responses of the teacher, all these relations are causal, thus the essential triangle is formed which makes communication about shared objects and events possible. but it is also this triangle that determines the content of the learner's words and thoughts when these become complex enough to deserve the term. the role of the teacher in determining the content of the learner's attitude is not just the 'determine' of causality. for in addition to being a cause of those thoughts, what makes the particular aspect of the cause of the learner's responses the aspect that gives them the content they have is the fact that this aspect of the cause is shared by the teacher and the learner. without such sharing, there would be no grounds for selecting one cause rather than another as the content-fixing cause. a noncommunicating creature may be seen by us as responding to an objective world; but we are not justified in attributing thoughts about our world (or any other) to it.

this brings me back to first person authority. i remarked at the start that first person authority has two closely related features: one is the fact that it yields knowledge that is not based on inference or evidence; the other is the asymmetry between how we know our own minds and how we know the minds of others. it is striking, for example, that burge's explanation of first person authority fails to account for the fact that the same thing, namely what is in one person's mind, can also be known by someone else, though in a very different way. the missing part of the explanation is filled in, 1 think, when we recognize the way interaction with other people partly determines the contents of mental states. knowledge of one's

own mind is personal. but what individuates that state at the same time makes it accessible to others, for the state is individuated by causal interplay among three elements: the thinker, others with whom he communicates, and an objective world they know they share.

## 14 three varieties of knowledge

i know, for the most part, what i think, want, and intend, and what my sensations are. in addition, i know a great deal about the world around me, the locations and sizes and causal properties of the objects in it. i also sometimes know what goes on in other people's minds. each of these three kinds of empirical knowledge has its distinctive characteristics. what i know about the contents of my own mind i generally know without appeal to evidence or investigation. there are exceptions, but the primacy of unmediated self-knowledge is attested by the fact that we distrust the exceptions until they can be reconciled with the unmediated. my knowledge of the world outside of myself, on the other hand, depends on the functioning of my sense organs, and this causal dependence on the senses makes my beliefs about the world of nature open to a sort of uncertainty that arises only rarely in the case of beliefs about our own states of mind. many of my simple perceptions of what is going on in the world are not based on further evidence; my perceptual beliefs are simply caused directly by the events and objects around me. but my knowledge of the propositional contents of other minds is never immediate in this sense; i would have no access to what others think and value if i could not note their behavior.

of course all three varieties of knowledge are concerned with aspects of the same reality; where they differ is in the mode of access to reality.

the relations among the three sorts of empirical knowledge, particularly questions of conceptual priority, have long headed the list of philosopher's epistemological concerns, and they are my subject here. many familiar approaches to the question how the three sorts of knowledge are related take self-knowledge as primary,

perhaps because of its directness and relative certainty, and then attempt to derive knowledge of the external world from it; as a final step, they try to base knowledge of other minds on observations of behavior. this is not, needless to say, the only direction the derivation can take: one may instead accept knowledge of the external world, at least in some of its manifestations, as basic, and try to relate or reduce the other forms of knowledge to it. the elaboration of such reductive proposals, and the demonstration of their failure, constitutes much of the history of philosophy from descartes to the present. if many philosophers have turned away from these problems in recent years, it is not because the problems are thought to have been solved, but because the problems seem intractable. there is also, of course, the wistful hope that the problems themselves are illusory.

this cannot be the case. there are compelling reasons for accepting the view that none of the three forms of knowledge is reducible to one or both of the others. here i give my own reasons for believing this; but i take the hopelessness of finding effective modes of reduction to be apparent from the almost universal rejection of standard reductionist programs. skepticism in various of its familiar guises is our grudging tribute to the apparent impossibility of unifying the three varieties of knowledge: one form of skepticism springs from the difficulty of accounting for our knowledge of the external world on the basis of our knowledge of our own minds; another recognizes that our knowledge of other minds cannot consist only in what we can observe from the outside. the intractability of the mindbody problem is another such tribute.

it is striking the extent to which philosophers, even those who have been skeptics about the possibility of justifying beliefs about the external world, have put aside these doubts when they have come to consider the problem of other minds; striking, since the latter problem can arise only if knowledge of behavior, and hence of the external world, is possible. holding the problems apart has the unfortunate effect of obscuring the fact that the two problems rest on a common assumption. the assumption is that the truth concerning what a person believes about the world is logically independent of the truth of those beliefs. this certainly seems to be the case, for surely the totality of a person's beliefs and subjective experiences is logically consistent with the falsity of any of those beliefs. so no amount of knowledge of the contents of one's own mind insures the truth of a belief about the external world. the logical independence of the mental works equally in the other direction: no amount of knowledge of the external world entails the truth about the workings of a mind. if there is a logical or epistemic barrier between the mind and nature, it not only prevents us from seeing out; it also blocks a view from outside in.

it is sometimes thought that if we separate the problem of knowing what is in a mind from the problem of knowing about anything whatever outside of ourselves, then the problem of knowledge of other minds is solved when we recognize that it is part of the concept of a mental state or event that certain forms of behavior, or other outward signs, count as evidence for the existence of that mental state or event. no doubt it is true that it is part of the concept of a mental state or event that behavior is evidence for it. what is unclear is how this answers the sceptic. for the fact that behavior is evidence for what is in a mind offers no explanation of the asymmetry between the indirect knowledge we have of other minds and the direct knowledge we have of our own mind. the proffered solution insists that behavioral evidence can suffice for the justified attribution of mental states to others, while it recognizes that such evidence is generally irrelevant to self-ascriptions of the same states. but if we are given no explanation of this striking asymmetry, we ought to conclude that there are really two kinds of concepts: mental concepts that apply to others, and mental concepts that apply to ourselves. if the mental states of others are known only through their behavioral and other outward manifestations, while this is not true of our own mental states, why should we think our own mental states are anything like those of others? we might also wonder why, if this answer to the problem of knowledge of other minds is satisfactory, we should not accept an analogous solution to the problem of our knowledge of the external world, yet it is widely recognized that this answer to general skepticism is unacceptable. do we distinguish between the problems because we suppose that while we have no access to the outside world except through experience, we nevertheless can intelligibly extrapolate to the experiences of others, since we have access to experience in our own case? but this supposition begs the question, since it assumes without argument that what we call the mental states of others are similar to what we identify as mental states in ourselves.

i have been rehearsing these problems and perplexities because i want, first of all, to stress the apparent oddity of the fact that we have

three irreducibly different varieties of empirical knowledge. we need an overall picture which not only accommodates all three modes of knowing, but makes sense of their relations to one another. without such a general picture we should be deeply puzzled that the same world is known to us in three such different ways. and, second, it is essential to appreciate the extent to which problems that have usually been taken one at a time are interrelated. there are three basic problems: how a mind can know the world of nature, how it is possible for one mind to know another, and how it is possible to know the contents of our own minds without resort to observation or evidence. it is a mistake, i shall urge, to suppose that these questions can be collapsed into two, or taken in isolation.

in trying to form a picture of the relations among the three kinds of knowledge we must do much more than show that they are mutually irreducible; we must see *why* they are irreducible. this in turn will involve bringing out the respective conceptual roles played by each of the forms of knowledge, and why each of these three sorts of knowledge is indispensable—why we could not get along without all of them. of course, if i am right that each of the three varieties of knowledge is indispensable, scepticism of the senses and scepticism about other minds must be dismissed. for the cartesian or humean skeptic about the external world holds that it is all too obvious that we can get along without knowledge of the world of nature—what we know of our own mind is self-sufficient, and may be all the knowledge we have. the skeptic about other minds is equally convinced that we can get along without knowledge of other minds-this must be possible if we are forever uncertain whether we have it.

it may seem at first that we could rather easily get along without a form of words to express our beliefs about the mental states of others or of ourselves. i think this is imaginable; but the issue with which i am concerned is primarily epistemic, not linguistic. the epistemic question is whether we could get along without knowledge of minds, both our own and those of others. i shall argue that we could not. what we could not do is get along without a way of expressing, and thus communicating, our thoughts about the natural world. but if we can do this, the transition to also being able verbally to attribute thoughts is relatively simple, and it would be astonishing if this step were not taken. with respect to our own thoughts, it is no more than the difference between saying assertively 'snow is white' and saying assertively 'i believe that snow is white'. the truth conditions of these assertions are not the same, but anyone who understands the first assertion knows the truth conditions of the second, even if he does not command a sentence with those truth conditions. this is because anyone who understands speech can recognize assertions, and knows that someone who makes an assertion represents himself as believing what he says. similarly, someone who says to jones that snow is white knows the truth conditions of 'jones believes that snow is white' (even if he does not know english nor has a way of expressing belief).

belief is a condition of knowledge. but to have a belief it is not enough to discriminate among aspects of the world, to behave in different ways in different circumstances; a snail or a periwinkle does this. having a belief demands in addition appreciating the contrast between true belief and false, between appearance and reality, mere seeming and being. we can, of course, say that a sunflower has made a mistake if it turns towards an artificial light as if it were the sun, but we do not suppose the sunflower can think it has made a mistake, and so we do not attribute a belief to the sunflower. someone who has a belief about the world—or anything else—must grasp the concept of objective truth, of what is the case independent of what he or she thinks. we must ask, therefore, after the source of the concept of truth.

wittgenstein put us on the track of the only possible answer to this question, whether or not his problem was as broad as ours, and whether or not he believed in answers to philosophical problems. the source of the concept of objective truth is interpersonal communication. thought depends on communication. this follows at once if we suppose that language is essential to thought and we agree with wittgenstein that there cannot be a private language.<sup>1</sup> the central argument against private languages is that, unless a language is shared, there is no way to distinguish between using the language correctly and using it incorrectly; only communication with another

<sup>&</sup>lt;sup>1</sup> of course there can be a private code based on a publicly acquired language. i have no idea how broadly wittgenstein intended his thesis about private languages to be interpreted: perhaps he intended his argument to apply only to those concepts which are necessarily private. but i, like saul kripke in *wittgenstein on rules and private language*, think the argument applies to language quite generally, and so (i would say) to prepositional thought, but while i accept the idea that communication is the source of objectivity, i do not think communication depends on speakers using the same words to express the same thoughts.
can supply an objective check. if only communication can provide a check on the correct use of words, only communication can supply a standard of objectivity in other domains, or so i shall argue. we have no grounds for crediting a creature with the distinction between what is thought to be the case and what is the case unless the creature has the standard provided by a shared language; and without this distinction there is nothing that can clearly be called thought.

in communication, what a speaker and the speaker's interpreter must share is an understanding of what the speaker means by what he says. how is this possible? it might help if we knew how language came into existence in the first place, or at least could give an account of how an individual learns his first language, given that others in his environment are already linguistically accomplished. failing such knowledge or account, what we can do instead is ask how a competent interpreter (one with adequate conceptual resources and a language of his own) might come to understand the speaker of an alien tongue. an answer to this question should reveal some important features of communication, and throw indirect light on what makes possible a first entry into language.

the intrepid interpreter, working without a bilingual trot, seeks to assign a propositional content to the utterances of a speaker. in effect he assigns a sentence of his own to each of the sentences of the speaker. to the extent that he gets things right, the interpreter's sentences provide the truth conditions of the speaker's sentences, and hence supply the basis for the interpretation of the speaker's utterances. the result can be thought of as a recursive characterization of truth, by the interpreter, of the sentences, and hence actual and potential utterances, of the speaker.

an interpreter cannot directly observe another person's propositional attitudes; beliefs, desires, and intentions, including the intentions which partly determine the meanings of utterances, are invisible to the naked eye. the interpreter can, however, attend to the outward manifestations of these attitudes, including utterances. since we are all able to discover from such manifestations what an agent thinks and means, there must be an intelligible relation between evidence and attitude. how do we bridge this gap? i know of only one way: an interpreter can perceive, often enough, that an agent has a certain sort of attitude towards an object or event the interpreter perceives. if the interpreter could in this way directly *individuate* the attitudes of someone else, the problem would be solved, but only by assuming the

#### three varieties of knowledge 211

interpreter is a mind reader. it would not beg the question, however, to assume the interpreter can detect one or more nonindividuating attitude. examples of the kind of special attitude i have in mind are: holding a sentence true at a time, wanting a sentence to be true, or preferring that one sentence rather than another be true. the assumption that we can detect such an attitude does not beg the question of how we endow the attitudes with content, since a relation such as holding true between a speaker and an utterance is an extensional relation which can be known to hold without knowing what the sentence means, these attitudes are nonindividuative, for though they are psychological in nature, they do not distinguish the different propositional contents different utterances express.

in *word and object* quine appealed to the nonindividuative attitude of prompted assent. since someone assents to an utterance, or holds a sentence true, in part because of what he believes and in part because of what the utterance or sentence means in his language, quine's problem was to separate out these two elements on the basis of evidence that combined their influence. if the separation succeeds, the result is a theory of both belief and meaning for the speaker, for it must yield an interpretation of the speaker's utterances, and if one knows both that the speaker assents to the utterance, and what it means in his mouth, one also knows what he believes.

the process of separating meaning and opinion invokes two key principles which must be applicable if a speaker is interpretable: the principle of coherence and the principle of correspondence. the principle of coherence prompts the interpreter to discover a degree of logical consistency in the thought of the speaker; the principle of correspondence prompts the interpreter to take the speaker to be responding to the same features of the world that he (the interpreter) would be responding to under similar circumstances. both principles can be (and have been) called principles of charity: one principle endows the speaker with a modicum of logic, the other endows him with a degree of what the interpreter takes to be true belief about the world. successful interpretation necessarily invests the person interpreted with basic rationality. it follows from the nature of correct interpretation that an interpersonal standard of consistency and correspondence to the facts applies to both the speaker and the speaker's interpreter, to their utterances and to their beliefs.

two questions now obtrude. the first is: why should an interpersonal standard be an objective standard, that is, why should what

people agree on be true? the second is: even if it is the case that communication assumes an objective standard of truth, why should this be the only way such a standard can be established?

here is a way of answering these questions. all creatures classify objects and aspects of the world in the sense that they treat some stimuli as more alike than others. the criterion of such classifying activity is similarity of response. evolution and subsequent learning no doubt explain these patterns of behavior. but from what point of view can these be called patterns? the criterion on the basis of which a creature can be said to be treating stimuli as similar, as belonging to a class, is the similarity of the creatures's responses to those stimuli; but what is the criterion of similarity of responses? this criterion cannot be derived from the creature's responses; it can only come from the responses of an observer to the responses of the creature. and it is only when an observer consciously correlates the responses of another creature with objects and events of the observer's world that there is any basis for saying the creature is responding to those objects or events rather than any other objects or events. as wouldbe interpreters of the verbal behavior of the speaker of an alien language, we group distinct verbal acts of the speaker together: 'mother', 'snow', 'table', when repeated as one-word sentences, sound similar if we are appropriately attuned, when we discover kinds of objects or events in the world that we can correlate with the utterances of a speaker, we are on the way to interpreting the simplest linguistic behavior.

if we are teaching someone a language, the situation becomes more complex, but more clearly interpersonal. what seems basic is this: an observer (or teacher) finds (or instills) a regularity in the verbal behavior of the informant (or learner) which he can correlate with events and objects in the environment. this much can take place without developed thought on the part of the observed, of course, but it is the necessary basis for attributing thoughts and meanings to the person observed. for until the triangle is completed connecting two creatures, and each creature with common features of the world, there can be no answer to the question whether a creature, in discriminating between stimuli, is discriminating between stimuli at the sensory surfaces or somewhere further out, or further in. without this sharing of reactions to common stimuli, thought and speech would have no particular content-that is, no content at all. it takes two points of view to give a location to the cause of a thought, and thus

to define its content. we may think of it as a form of triangulation: each of two people is reacting differentially to sensory stimuli streaming in from a certain direction. projecting the incoming lines outward, the common cause is at their intersection. if the two people now note each other's reactions (in the case of language, verbal reactions), each can correlate these observed reactions with his or her stimuli from the world. a common cause has been determined. the triangle which gives content to thought and speech is complete. but it takes two to triangulate.

until a base line has been established by communication with someone else, there is no point in saying one's own thoughts or words have a propositional content. if this is so, then it is clear that knowledge of another mind is essential to all thought and all knowledge. knowledge of another mind is possible, however, only if one has knowledge of the world, for the triangulation which is essential to thought requires that those in communication recognize that they occupy positions in a shared world. so knowledge of other minds and knowledge of the world are mutually dependent; neither is possible without the other. ayer was surely right when he said, 'it is only with the use of language that truth and error, certainty and uncertainty, come fully upon the scene'.<sup>2</sup>

knowledge of the propositional contents of our own minds is not possible without the other forms of knowledge since there is no propositional thought without communication. it is also the case that we are not in a position to attribute thoughts to others unless we know what we think since attributing thoughts to others is a matter of matching the verbal and other behavior of others to our own propositions or meaningful sentences. knowledge of our own minds and knowledge of the minds of others are thus mutually dependent.

it should now be clear what insures that our view of the world is, in its plainest features, largely correct. the reason is that the stimuli that cause our most basic verbal responses also determine what those verbal responses mean, and the content of the beliefs that accompany them. the nature of interpretation guarantees both that a large number of our simplest perceptual beliefs are true, and that the nature of these beliefs is known to others. of course many beliefs are given content by their relations to further beliefs, or are caused by misleading sensations; any particular belief or set of beliefs about the world

<sup>2</sup> a. j.ayer. the problem of knowledge. 54.

around us may be false. what cannot be the case is that our general picture of the world and our place in it is mistaken, for it is this picture which informs the rest of our beliefs and makes them intelligible, whether they be true or false.

the assumption that the truth of what we believe is logically independent of what we believe is revealed as ambiguous. any particular belief may indeed be false; but enough in the framework and fabric of our beliefs must be true to give content to the rest. the conceptual connections between our knowledge of our own minds and our knowledge of the world of nature are not definitional but holistic. the same is true of the conceptual connections between our knowledge of behavior and our knowledge of other minds.

there are, then, no 'barriers', logical or epistemic, between the three varieties of knowledge. on the other hand, the very way in which each depends on the others shows why none can be eliminated, or reduced to the others.

as noted above, we may think of an interpreter who aims to understand a speaker as matching up sentences of his own with the utterances and states of mind of the speaker. the totality of evidence available to the interpreter determines no unique theory of truth for a given speaker, not just because actually available evidence is finite while the theory has an infinity of testable consequences, but because all possible evidence cannot limit acceptable theories to one. given the richness of the structure represented by the set of one's own sentences, and the nature of the connections between the members of this set and the world, we should not be surprised if there are many ways of assigning our own sentences to the sentences and thoughts of someone else that capture everything of significance.

the situation is analogous to the measurement of weight or temperature by assigning numbers to objects. even supposing there are no errors of measurement, and that all possible observations have been made, an assignment of numbers to objects that correctly registers their weights is not unique: given one such assignment, another can be produced by multiplying all the numbers by any positive constant. in the case of ordinary temperature (not absolute temperature), any correct assignment of numbers can be converted to another by a linear transformation. because there are many different but equally acceptable ways of interpreting an agent we may say, if we please, that interpretation or translation is indeterminate, or that there is no fact of the matter as to what someone means by his or her words. in the same vein, we could speak of the indeterminacy of weight or temperature. but we normally accentuate the positive by being clear about what is invariant from one assignment of numbers to another, for it is what is invariant that is empirically significant. the invariant *is* the fact of the matter. we can afford to look at translation and the content of mental states in the same light.<sup>3</sup>

i once thought that the indeterminacy of translation supplied a reason for supposing there are no strict laws connecting mental and physical concepts, and so supported the claim that mental concepts are not even nomologically reducible to physical concepts. i was wrong: indeterminacy turns up in both domains. but one source of indeterminacy in the case of the mental is that the line between empirical truth and truth due to meaning cannot in general be clearly defined on behavioral grounds; and behavioral grounds are all we have for determining what speakers mean. it is here that the irreducible difference between mental concepts and physical concepts begins to emerge: the former, at least insofar as they are intentional in nature, require the interpreter to consider how best to render the creature being interpreted intelligible, that is, as a creature endowed with reason, as a consequence, an interpreter must separate meaning from opinion partly on normative grounds by deciding what, from his point of view, maximizes intelligibility. in this endeavor the interpreter has, of course, no other standards of rationality to fall back on than his own, when we try to understand the world as physicists, we necessarily employ our own norms, but we do not aim to discover rationality in the phenomena.

how does the normative element in mental concepts prevent their reduction to physical concepts? perhaps it is obvious that definitional reduction is out of the question; but why can't there be laws—strict laws—that connect each mental event or state with events or states described in the vocabulary of an advanced physics? when writing about this twenty years ago i said, in effect, that one can hope for strict connecting laws only when the concepts connected by the laws are based on criteria of the same sort, and so a strict law could not combine normative with nonnormative concepts.<sup>4</sup> this answer still seems to me right as far as it goes, but

<sup>&</sup>lt;sup>3</sup> here i accept quine's thesis of the indeterminacy of translation, and extend it to the interpretation of thought generally. the analogy with measurement is my own.

in 'mental events', essay 11 in essays on actions and events.

it has understandably been found inconclusive by critics. i now want to add some further considerations.

one further consideration is this: strict laws do not employ causal concepts, while most, if not all, mental concepts are irreducibly causal. an action, for example, must be intentional under some description, but an action is intentional only if it is caused by mental factors such as beliefs and desires. beliefs and desires are identified in part by the sorts of action they are prone to cause, given the right conditions. many of the concepts that feature in commonsense explanations are causal in this way. an accident was caused by the fact that the road was slippery; something is slippery if it causes appropriate objects to slip under appropriate circumstances. we explain why the wing of an airplane does not break when it bends by noting that it is made of elastic materials; a material is elastic if there is something about it that causes it, under appropriate conditions, to return to its original shape after deformation. such explanations do not lend themselves to precision for two reasons: we cannot spell out in detail when the circumstances are appropriate, and the appeal to causality finesses part of what a full-scale explanation would make manifest. descriptions of objects, states, and events that are needed to instantiate strict, exceptionless laws do not contain causal concepts (which is not to say that laws which contain only noncausal concepts are not causal laws).

in the case of causal properties like elasticity, slipperiness, malleability, or solubility, we tend to think, rightly or wrongly, that what they leave unexplained can be (or already has been) explained by the advance of science. we would not be changing the subject if we were to drop the concept of elasticity in favor of a specification of the microstructure of the materials in the airplane wing that cause it to return to its original shape when exposed to certain forces. mental concepts and explanations are not like this. they appeal to causality because they are designed, like the concept of causality itself, to single out from the totality of circumstances which conspire to cause a given event just those factors that satisfy some particular explanatory interest. when we want to explain an action, for exam ple, we want to know the agent's reasons, so we can see for ourselves what it was about the action that appealed to the agent. but it would be foolish to suppose that there are strict laws that stipulate that whenever an agent has certain reasons he will perform a given action.

#### three varieties of knowledge 217

the normative and the causal properties of mental concepts are related. if we were to drop the normative aspect from psychological explanations, they would no longer serve the purposes they do. we have such a keen interest in the subject's reasons for acting and for his or her changes of belief that we are willing to settle for explanations that cannot be made to fit perfectly with the laws of physics. physics, on the other hand, has as an aim laws that are as complete and precise as we can make them; a different aim. the causal element in mental concepts helps make up for the precision they lack; it is part of the concept of an intentional action that it is caused and explained by beliefs and desires; it is part of the concept of a belief or a desire that it tends to cause, and so explain, actions of certain sorts.

much of what i have said about what distinguishes mental concepts from the concepts of a developed physics could also be said to distinguish the concepts of many of the special sciences such as biology, geology, and meteorology. so even if i am right that the normative and causal character of mental concepts divides them definitionally and nomologically from the concepts of a developed physics, it may seem that there must be something more basic or foundational that accounts for this division. i think there is.

knowledge of the contents of our own minds must, in most cases, be trivial. the reason is that, apart from special cases, the problem of interpretation cannot arise. when i am asked about the prepositional contents of my mind, i must use my own sentences. the answer is usually absurdly obvious: my sentence 'snow is white', like my thought that snow is white, is true if and only if snow is white. my knowledge of the contents of another mind is possible, i have argued, only in the context of a generally correct, and shared, view of the world. but such knowledge differs from the knowledge i have of my own mind since it is necessarily indirect in that it depends, among other things, on observed correlations between the speech and other behavior of the person, and of events in our communal environment.

the fundamental difference between my knowledge of another mind and of the shared physical world has a different source. communication, and the knowledge of other minds that it presupposes, is the basis of our concept of objectivity, our recognition of a distinction between false and true belief. there is no going outside this standard to check whether we have things right, any more than we can check whether the platinum-iridium standard kept at the

international bureau of weights and standards in sevres, france, weighs a kilogram. (this comparison was valid when the standard in sevres defined the kilogram.) we can, of course, turn to a third party and a fourth to broaden and secure the interpersonal standard of the real, but this leads not to something intrinsically different, just to more of the same.

i spoke before of an analogy between how we assign numbers to keep track of the relations among objects with respect to temperature or weight and how we use our own sentences to identify the contents of the thoughts and utterances of others. but the analogy is imperfect: the nature of the scaling device differs in the two cases. we depend on our linguistic interactions with others to yield agreement on the properties of numbers and the sort of structures in nature that allow us to represent those structures in the numbers. we cannot in the same way agree on the structure of the sentences or thoughts we use to chart the thoughts and meanings of others, for the attempt to reach such an agreement simply sends us back to the very process of interpretation on which all agreement depends.

it is here, i suggest, that we come to the ultimate springs of the difference between understanding minds and understanding the world as physical. a community of minds is the basis of knowledge; it provides the measure of all things. it makes no sense to question the adequacy of this measure, or to seek a more ultimate standard.

we have dwelt at length on the inescapability of the objective aspect of all thought. what remains of the subjective aspect? clearly we have not obliterated the difference between self-knowledge and knowledge of other minds: the first remains direct and the second indirect. objectivity itself we have traced to the intersections of points of view—for each person, the relation between his own reactions to the world and the reactions of others. these differences are real. our thoughts are 'inner' and 'subjective' in that we know what they are in a way no one else can. but though possession of a thought is necessarily individual, its content is not. the thoughts we form and entertain are located conceptually in the world we inhabit, and know we inhabit, with others. even our thoughts about our own mental states occupy the same conceptual space and are located on the same public map.

the philosophical conception of subjectivity is burdened with a history and a set of assumptions about the nature of mind and meaning that sever the meaning of an utterance or the content of a thought

#### three varieties of knowledge 219

from questions about external reality, 'my' world from the world as it appears to others. this popular conception holds that the subjective is prior to the objective, that there is a subjective world prior to knowledge of external reality. it is evident that the picture of thought and meaning i have sketched here leaves no room for such priority since it predicates self-knowledge on knowledge of other minds and of the world. the objective and the intersubjective are thus essential to anything we can call subjectivity, and constitute the context in which it takes form. collingwood put it succinctly:

the child's discovery of itself as a person is also its discovery of itself as a member of a world of persons . . . the discovery of myself as a person is the discovery that i can speak, and am thus a *persona* or speaker; in speaking i am both speaker and hearer; and since the discovery of myself as a person is also the discovery of other persons around me, it is the discovery of speakers and hearers other than myself.<sup>5</sup>

it may seem that if sharing a general view of the world is a condition of thought, the differences in intellectual and imaginative character among minds and cultures will be lost to sight. if i have given this impression, it is because i have wanted to concentrate on what seems to me primary, and so apt to go unnoticed: the necessary degree of communality essential to understanding another individual, and the extent to which such understanding provides the foundation of the concept of truth and reality upon which all thought depends. but i do not want to suggest that we cannot understand those with whom we differ on vast tracts of physical and moral opinion. it is also the case that understanding is a matter of degree: others may know things we do not, or even perhaps cannot. what is certain is that the clarity and effectiveness of our concepts grows with the growth of our understanding of others. there are no definite limits to how far dialogue can or will take us.

some philosophers worry that if all our knowledge, at least our propositional knowledge, is objective, we will lose touch with an essential aspect of reality: our personal, private outlook. i think this worry is groundless. if i am right, our propositional knowledge has its basis not in the impersonal but in the interpersonal. thus, when we look at the natural world we share with others, we do not lose contact with ourselves, but rather acknowledge membership in a society of minds. if i did not know what others think, i would have

<sup>5</sup> r. g. collingwood, *the principles of art*, 248.

no thoughts of my own and so would not know what i think. if i did not know what 1 think, i would lack the ability to gauge the thoughts of others. gauging the thoughts of others requires that i live in the same world with them, sharing many reactions to its major features, including its values. so there is no danger that in viewing the world objectively we will lose touch with ourselves. the three sorts of knowledge form a tripod: if any leg were lost, no part would stand.

# contents list of volumes of essays by donald davidson

volume 1

## essays on actions and events

intention and action

- 30. actions, reasons, and causes (1963)
- 31. how is weakness of the will possible? (1969)

ь

- 32. agency (1971)
- 33. freedom to act (1973)
- 34. intending (1978)

#### event and cause

- 35. the logical form of action sentences (1967)
- 36. causal relations (1967)
- 37. the individuation of events (1969)
- 38. events as particulars (1970)
- 10. eternal vs. ephemeral events (1971)

#### philosophy of psychology

- 39. mental events (1970)
- 40. psychology as philosophy (1974)
- 41. the material mind (1973)
- 42. hempel on explaining action (1976)
- 43. hume's cognitive theory of pride (1976)

#### appendices

- a. adverbs of action (1985)
- b. reply to quine on events (1985)

#### volume 2

### inquiries into truth mid interpretation

#### truth and meaning

- 44. theories of meaning and learnable languages (1965)
- 45. truth and meaning (1967)
- 46. true to the facts (1969)
- 47. semantics for natural languages (1970)
- 48. in defence of convention t (1973)

#### applications

- 49. quotation (1979)
- 50. on saying that (1968)
- 51. moods and performances (1979)

#### radical interpretation

- 9. radical interpretation (1973)
- 52. belief and the basis of meaning (1974)
- 53. thought and talk (1975)
- 54. reply to foster (1976)

#### language and reality

- 55. on the very idea of a conceptual scheme (1974)
- 56. the method of truth in metaphysics (1977)
- 57. reality without reference (1977)
- 58. the inscrutability of reference (1979)

#### limits of the literal

- 59. what metaphors mean (1978)
- 60. communication and convention (1982)

appendix to essay 10 replies to quine and lewis (1974)

,

volume 3

### subjective, intersubjective, objective

subjective

- 61. first person authority (1984)
- 62. knowing one's own mind (1987)
- 63. the myth of the subjective (1988)
- 64. what is present to the mind (1989)
- 65. indeterminism and antirealism (1997)
- 66. the irreducibility of the concept of the self (1998)

intersubjective

*1.* rational animals (1982) 8. the second person (1992) 9 the emergence of thought (1997)

#### objective

 a coherence theory (1983) afterthoughts (1987) 1 1. empirical content (1982) 1 2. epistemology and truth (1988)
epistemology externalized (1990)
three varieties of knowledge (1991)

volume 4

## problems of rationality

#### rationality

- 69. the problem of objectivity (1995)
- 70. expressing evaluations (1984)
- 71. the objectivity of value (1995)
- 72. judging interpersonal values (1996)?'
- 73. objectivity and practical reason

#### problems and proposals

- 74. turing's test (1990)
- 75. representation and interpretation (1990)
- 76. problems of explaining action (1987)
- 77. could there be a theory of rationality? (1995)
- 78. what thought requires
- 79. the unified theory of thought and action (1980)

#### irrationality

- 80. paradoxes of irrationality (1974)
- 81. incoherence and irrationality (1985)
- 82. deception and division (1986)
- 83. who is fooled? (1997)

volume 5

## truth, language and history

truth

- 84. truth rehabilitated (2000)
- 85. the folly of trying to define truth (1996) replies to rorty, stroud, and mcdowell
- 86. meaning, truth, and evidence (1990)
- 87. pursuit of the concept of truth (1995)
- 88. what is quine's view of truth? (1994)

#### language

- 89. a nice derangement of epitaphs (1986)
- 90. the social aspect of language (1994)
- 91. seeing through language (1997)
- 92. joyce and hurnpty dumpty (1989)
- 93. the third man (1992)
- 94. locating literary language (1993)

#### historical thoughts

- 95. plato's philosopher (1985)
- 96. the socratic concept of truth (1992)
- 97. dialectic and dialogue (1994)
- 98. gadamer and plato's philebus (1997)
- 99. aristotle's action
- 100.method and metaphysics (1993)

## bibliographical references

AGASSI, JOSEPH, science in flux (dordrecht: reidel, 1975).

ALSTON, WILLIAM, 'varieties of privileged access', *american philosophical quarterly*, 9(1971). AVRAMIDES, ANITA, 'davidson and the new sceptical problem', in

u. zeglen (ed.), *donald davidson: truth, meaning and knowledge* (london: routledge, 1999).

AYER, a. j., *the problem of knowledge* (london: macmillan, 1956). AYER, a. j. (ed.), *logical positivism* (new york: free press, 1959). AYER, a. j., 'privacy', in *the concept of a person and other essays* (new

york: macmillan, 1963).

BENNETT, JONATHAN, *linguistic behavior* (cambridge university press, 1976). BLACK, MAX, the semantic definition of truth', *analysis*, 8 (1948). BRENTANO, FRANZ, *psychology from an empirical standpoint*, ed. o. kraus,

trans. a. rancurello, d. b. terrell, and l. l. mcalister (london: humanities press, 1973). BURGE, TYLER, 'individualism and the

mental', in p. a. french,

t. e. uehling, and h. k. wettstein (eds.), *midwest studies in philosophy* (minneapolis: university of minnesota press, 1979). BURGE, TYLER, 'two thought experiments reviewed', *notre dame journal* 

of formal logic, 23 (1982). BURGE, TYLER, 'other bodies', in a. woodfield (ed.), thought and object

(oxford university press, 1982). BURGE, TYLER, 'individualism and psychology', *philosophical review*, 95

(1986). BURGE, TYLER, 'cartesian error and the objectivity of perception', in

p. petit and j. mcdowell (eds.), *subject, thought, and context* (oxford university press, 1986). BURGE, TYLER, 'intellectual norms and foundations of mind', *journal of* 

*philosophy*, 83(1986). BURGE, TYLER, 'perceptual individualism and authoritative self-knowledge'.

in r. grimm and d. merrill (eds.). *contents of thought* (tucson: university of arizona press, 1988). BURGE, TYLER, 'individualism and self-knowledge', *journal of* 

philosophy, 85 (1988).

#### 228 bibliographical references

CARNAP, RUDOLF, 'psychology in physical language', *erkenntnis*, 3 (1932-3). CHAN, WAI PANG, FREDERICK PRETE, and MICHAEL h. DICKINSON,

'visual

input to the efferent control system of a fly's "gyroscope" ', *science*, 280 (10 apr. 1998). chomsky, noam, *knowledge of language: its nature*, *origin and use* 

(new york: praeger, 1986). CHOMSKY, NOAM, language and problems of knowledge (cambridge,

mass.: mit press, 1988). CHURCHLAND, PAUL, scientific realism and the plasticity of mind

(cambridge university press, 1979). COLLINGWOOD, r. g., the principles of art (oxford university press,

1938). DAVIDSON, DONALD, essays on actions and events (oxford university

press, 1980). DAVIDSON, DONALD, *inquiries into truth and interpretation* (oxford

university press, 1980). DAVIDSON, DONALD, 'a nice derangement of epitaphs', in r. grandy and

r. warner, (eds.), *(philosophical grounds of rationality* (oxford university press, 1986). DAVIDSON, DONALD, 'meaning, truth and evidence', in r. barrett and r.

gibson (eds.), perspectives on quine (oxford: blackwell, 1990).

DAVIDSON, DONALD, 'the structure and content of truth', *journal of philosophy*, 87(1990). DAVIDSON, DONALD, 'three varieties of knowledge', in

a. p. griffiths (ed.),

*a. j. ayer memorial essays* (cambridge university press, 1991). DAVIDSON, DONALD, ' pursuit of the concept of truth', in p. leonard) and

m. santambrogio (eds.), *on quine* (cambridge university press, 1995). DAVIDSON, DONALD, 'replies', *critica*, 30 (1998). DENNETT, DANIEL, 'beyond belief, in a. woodfield (ed.), *thought and* 

object (oxford university press, 1982).

DENNETT, DANIEL, 'real patterns', *journal of philosophy*, 88 (1991). DUMMETT, MICHAEL, *truth and other enigmas* (london: duckworth,

1978). DUMMETT, MICHAEL, the interpretation of frege's philosophy

(london:

duckworth, 1981).

EVANS, GARETH, *the varieties of reference* (oxford university press, 1982). FODOR, JERRY, 'methodological solipsism considered as a research

strategy in cognitive psychology', *behavioral and brain sciences*, 3 (1980). FODOR, JERRY, 'cognitive science and the twin earth problem', *notre* 

*dame journal of formal logic*, 23 (1982). FODOR, JERRY, *psychosemantics* (cambridge, mass.: mit press, 1987).

FODOR, JERRY, *a theory of content* (cambridge, mass.: mit press, 1990). f0llesdal, DAGKINN, 'knowledge, identity, and existence', *theoria*, 33

(1967). GOODMAN, NELSON, *fact, fiction, and forecast,* 4th edn. (cambridge,

mass.: harvard university press, 1983). GUHRAUER, g., gottfried wilhelm freiherr von leibniz:. eine biographic

(hildesheim: gb. olms, 1966). HAHN, LEWIS e. (ed.), the philosophy of donald davidson, library of

living philosophers (chicago: open court, 1999). HAHN, LEWIS e., and p. a. SCHILPP (eds.), *the philosophy of w. v. quine*,

library of living philosophers (chicago: open court, 1986). HEMPEL, CARL, 'on the logical positivist's theory of truth', *analysis*, 2

(1935). HEMPEL, CARL, 'the empiricist criterion of meaning', *revue international* 

*de philosophic*, 4(1950). HEMPEL, CARL, *aspects of scientific explanation* (new york: free press,

1965).

HENRICH, DIETER (ed.), *kant oder hegel?* (stuttgart: klett-kotta, 1983). JEFFREY, RICHARD, *the logic of decision* (university of chicago press, 1965). KAPLAN, DAVID, 'quantifying in', *synthese*, 19 (1968). KRIPKE, SAUL, 'naming and necessity', in d. davidson and g. harman

(eds.), *semantics of natural language* (dordrecht: reidel, 1972). KRIPKE, SAUL, *wittgenstein on rules and private language* (cambridge:

blackwell, 1982). LEPORE, ERNIE (ed), truth and interpretation: perspectives on the

philosophy of donald davidson (cambridge: blackwell, 1986). LEWIS, c. i., mind and the world order (new york: scribner's, 1960). LEWIS, DAVID,

'languages and language', in k. gunderson (ed.), *language, mind and knowledge* (minneapolis: university of minnesota press, 1975).

MALACHOWSKI, ALAN (ed.), *reading rorty* (oxford: blackwell, 1990). MALCOLM, NORMAN, 'thoughtless brutes', *proceedings and addresses of* 

the american philosophical association, 46 (1972-3). MARGALIT, AVISHAI (ed.). meaning and use (dordrecht: reidel, 1979). MATES, BENSON, the philosoph of leibniz: metaph sics and language

(oxford university press, 1986). NAGEL, THOMAS, the objective self, in c. genet and s. shoemaker (eds.),

*mind and knowledge* (oxford university press, 1983). NEURATH, OTTO, 'protocol sentences', *erkenntnis*, 3 (1932-3). PUTNAM, HILARY, 'the meaning of "meaning" ', in *philosophical papers*,

ii: *mind, language, and reality* (cambridge university press. 1975). PUTNAM, HILARY, *meaning and the moral sciences* (london: routledge & kegan paul, 1978).

#### 230 bibliographical references

PUTNAM, HILARY, *realism and reason* (cambridge university press, 1983). QUINE, W. V., 'quantifiers and prepositional attitudes', *journal of philosophy*, *53* (1956).

QUINE, W. V., word and object (cambridge, mass.: mit press, i960). QUINE, W. V., ontological relativity and other essays (new york:

columbia university press, 1969).

QUINE, W. V., *the roots of reference* (la salle, iii.: open court, 1974). QUINE, W. V., the nature of natural knowledge', in s. guttenplan (ed.),

*mind and language* (oxford university press, 1975). QUINE, W. V., 'on empirically equivalent systems of the world',

*erkenntnis*, 9(915). QUINE, W. v, 'intensions revisited', in *theories and things* (cambridge,

mass.: harvard university press, 1981). QUINE, W. V., *theories and things* (cambridge, mass.: harvard university

press, 1981). QUINE, w. v., 'reply to stroud', in p. a. french, t. e. uehling, and

h. k. wettstein (eds.), *midwest studies in philosophy* (minneapolis: university of minnesota press, 1981). QUINE, w. v., 'events and miffordian', in a langer and h. malauch lin

reification', in e. lepore and b. mclaughlin

(eds.), actions and events: perspectives on the philosophy of donald davidson (cambridge: blackwell, 1985). QUINE, W. V., pursuit of truth, rev. edn. (cambridge, mass.: harvard

university press, 1992). QUINE, W. V., *from stimulus to science*, (cambridge, mass.: harvard

university press, 1995).  $\ensuremath{\text{QUINE}}$  , w. v., and  $\ensuremath{\text{Nelson}}$  goodman, 'steps toward a constructive

nominalism', *journal of symbolic logic*, 12 (1947), 97-122. RAMSEY, FRANK, 'truth and probability', in *philosophical papers:* 

*f. p. ramsey*, ed. d. h. mellor (cambridge university press, 1990). RORTY, RICHARD, 'incorrigibility as the mark of the mental', *journal of* 

philosophy, 67 (1970). RORTY, RICHARD, philosophy and the mirror of nature (princeton

university press, 1979). RORTY, RICHARD, *consequences of pragmatism* (minneapolis: university of

minnesota press, 1982). RORTY, RICHARD, 'pragmatism, davidson and truth', in e. lepore (ed.),

truth and interpretation: perspectives on the philosophy of donald davidson (oxford: blackwell, 1986).

RYLE, GILBERT, the concept of mind (new york: barnes & noble, 1949).

SCHIFFER, STEPHEN, remnants of meaning (cambridge, mass.: mit press,

1987). SCHLICK, MORITZ, 'the turning point in philosophy', erkenntnis,

*l* (1930-1).

SCHLICK, MORITZ, 'positivism and realism', *erkenntnis*, 3 (1932-3). SCHLICK, MORITZ, the foundation of knowledge', *erkenntnis*, 4 (1934). SCHLICK, MORITZ, 'meaning and verification', *philosophical review*,

4(1936).

SEARLE, JOHN, *intentionality* (cambridge university press, 1983). SEARLE, JOHN, 'indeterminacy and the first person', *journal of philosophy*,

84(1987). SELLARS, WILFRED, 'empiricism and the philosophy of mind', in h. feigl

and m. scriven (eds.), *minnesota studies in the philosophy of science* (minneapolis: university of minnesota press, 1956). SHOEMAKER, SIDNEY, *self-knowledge and self-identity* (ithaca, ny: cornell

university press, 1963). STICK, STEPHEN, 'autonomous psychology and the belief-desire thesis',

*monist*, *61* (1978). STICK, STEPHEN, *from folk psychology to cognitive science* (cambridge,

mass.: mit press, 1983).

STRAWSON, PETER, *individuals* (london: methuen, 1959). STRAWSON, PETER, truth', in *logico-linguistic papers* (london: methuen,

1971). TARSKI, ALFRED, the concept of truth in formalized languages', in

j. h. woodger (ed.), *logic, semantics, metamathematics* (oxford university press, 1956).

WALLAS, GRAHAM, *the art of thought* (new york: harcourt brace, 1926). WEISS, DONALD, 'professor malcolm on animal intelligence',

*philosophical review*, 84 (1975). WITTGENSTEIN, LUDWIG, *philosophical investigations* (new york:

macmmillan, 1953). WOODFIELD, ANDREW (ed.), *thought and object* (oxford university press,

1982). WRIGHT, CRISPIN, 'kripke's account of the argument against private

language', journal of philosophy, 81 (1984), 759-78.