New Horizons in the Study of Language and Mind

Noam Chomsky
Massachusetts Institute of Technology

Language & Interpretation

CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK www.cup.cam.ac.uk
40 West 20th Street, New York, NY 10011–4211, USA www.cup.org
10 Stamford Road, Oakleigh, Melbourne 3166, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain

Aviva Chomsky and Eric F. Menoya as Trustees of the Diane Chomsky Irrevocable Trust
© Foreword: Neil Smith 2000

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2000

Printed in the United Kingdom at the University Press, Cambridge

Typeset in 10/12 pt Plantin [GC]

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Chomsky, Noam.
New horizons in the study of language and mind / Noam Chomsky.
p. cm.
Includes bibliographical references and index.
ISBN 0 521 65147 6 (hardback) – ISBN 0 521 65822 5 (paperback)
1. Language and languages – Philosophy. 2. Philosophy of mind.
I. Title
P106.C524 2000
401–dc21 99–36753 CIP

ISBN 0 521 65147 6 hardback
ISBN 0 521 65822 5 paperback

Contents

3 Language and interpretation: philosophical reflections and empirical inquiry

In the philosophical literature of the past 40 years, there have been several influential currents that seem to me problematic in important, even essential respects. I have in mind, in the first place, approaches that take as their point of departure certain conceptions of how language is studied, or should be studied, by the empirical scientist – or the "field linguist," to use the terms of Quine's familiar paradigm. One can include here Quine, Donald Davidson, and others who have moved towards a form of pragmatism and "naturalized epistemology," incorporating questions thought to be of philosophical significance within their conception of empirical science, but also others who adopt a different starting point: Michael Dummett, and many of those influenced by Wittgenstein and ordinary language philosophy, for example.

To illustrate the flavor of these ideas, take some comments of Richard Rorty in Lepore (1986) on Davidson. He writes that "Davidson is surely right that Quine 'saved philosophy of language as a serious subject' by getting rid of the analytic–synthetic distinction. Quine's best argument for doing so was that the distinction is of no use to the field linguist" (Rorty 1986: 339).

As for the "field linguist," all that he "has to go on is his observation of the way in which linguistic is aligned with non-linguistic behavior in the course of the native's interaction with his environment, an interaction which [the linguist] takes to be guided by rules of action . . . ," specifically, the "regulative principle" that "most of the native's rules are the same as ours, which is to say that most of them are true" (p. 340; "rules" here apparently referring to beliefs). We need not be concerned about "a conceptual scheme, a way of viewing things, a perspective (or . . . a language, or a cultural tradition), [because] the field linguist does not need them, [so] therefore philosophy does not need them either" (p. 344). Quine and Davidson agree that "a theory of meaning for a language is what comes out of empirical research into linguistic behavior," when this is properly pursued, in accord with the doctrines of "holism and behaviorism" (p. 352).

This line of thought, Rorty continues, leads to a form of pragmatism that he espouses and attributes to James and Dewey, including crucially the denial of any relations of "'being made true' which hold between beliefs and the world." Rather, "We understand all there is to know about the relation of beliefs to the world when we understand their causal relations with the world" (p. 335).

Putting aside the conclusions that Rorty reaches,[1] consider his assumptions. If the best argument for dispensing with the analytic–synthetic distinction is that it is of no use to the field linguist, then virtually everyone who actually works in descriptive semantics, or ever has, must be seriously in error, since such work is shot through with assumptions about connections of meaning, which will (in particular) induce examples of the analytic-synthetic distinction. One would be hard put to find studies of language that do not assign structures and describe the meaning of *kill*, *so*, etc., in such a way that there is a qualitative distinction – determined by the language itself – between the sentences "John killed Bill, so Bill is dead," and "John killed Bill, so John is dead." Or, to take another case, it would be difficult to find a study of referential dependence in natural language that does not conclude that the language itself determines that the relation holds between *Mary* and *herself* in (1), but not when the same expression is embedded in the context "I wonder who –," yielding (2).

(1) Mary expects to feed herself.

(2) I wonder who Mary expects to feed herself.

Such syntactic–semantic properties will induce cases of the analytic–synthetic distinction; thus they will yield a distinction between "Mary expects to feed herself, so Mary expects to feed *Mary*" (analytic, with the three occurrences of *Mary* taken to be coreferential), and "I wonder who Mary expects to feed herself, so I wonder who Mary expects to feed *Mary*" (not analytic, under the same interpretation). But what Quine is alleged to have demonstrated goes beyond the matter of analyticity, reaching the conclusion that there are no semantic connections that can be attributed to the language faculty itself as distinct from our general systems of belief; elsewhere, Rorty takes this to be one of the two fundamental discoveries that undermine a traditional world picture.

As is well known, Quine and others have offered their own account of these distinctions. I return to these proposals, and how they might be evaluated in accordance with the canons of inquiry of the natural sciences, but merely note here that reference to "the field linguist" can surely not be understood as reference to those who actually do linguistic work. Rather, it has a normative character, referring to the way such

work ought to be done, keeping to the conditions of "holism and behaviorism" legislated by the philosopher, but not followed in practice by the errant scientist. While it might turn out on investigation that this stance is justifiable, those with an appreciation of the history of the discipline might be pardoned some initial skepticism.

To select another example to illustrate the flavor of these discussions, consider Dummett's argument in the same volume (Dummett 1986) that the "fundamental sense" in which we must understand the concept of language is the sense in which Dutch and German are different languages (he gives a different example, but the point is the same), each of them a particular social practice "in which people engage," a practice that "is learned from others and is constituted by rules which it is part of social custom to follow" (p. 473). Thus Dutch and German exist in this "fundamental sense," "independently of any particular speakers"; every individual speaker "has" such a language, but typically has only a "partial, and partially erroneous, grasp of the language." The intended import of Dummett's proposal is far-reaching. He is telling us what notion of "language" is essential for philosophical purposes, for the theory of meaning in particular; and also, as he makes clear, it is this concept of language that is in his view required for explaining the use of language, specifically, for understanding "what long-range theory someone brings to a first linguistic encounter with another." It is, therefore, a proposal that bears on the empirical study of language, of people, of what they know and what they do. Perhaps he means to allow that linguists may follow some different course for their special concerns, but clearly these proposals bear on the proper practice in empirical inquiry into language and its use.

Here the paradoxical flavor is of a somewhat different order. It lies in the conflict between Dummett's proposal and the commonplace assumption in empirical practice that there is no useful general sense in which we can characterize "language" so that Dutch and German are two distinct "languages," which people know only "partially" and "erroneously." This is so whether we are studying language structure, psycholinguistics, language change, typology, problems of communication, or whatever. People who live near the Dutch border can communicate quite well with those living on the German side, but they speak different languages in accordance with the sense of the term that Dummett argues is "fundamental"; and those on the German side of the border, with their "partial knowledge" of the "language German", may understand nothing spoken by people living in some other region, who "have" a different "partial knowledge" of the "language German" in Dummett's sense. It is for such reasons as these that no such concept

plays any role in empirical inquiry into language or psychology. Such terms as "English" and "Japanese" are used for general expository discourse, but with the understanding that their common-sense usage, which Dummett rather uncritically adopts, is to be abandoned when we turn to actual study of language, behavior, and communication.[2] If Dummett's concept is indeed fundamental for empirical inquiry and for philosophical purposes, then either philosophy, or the empirical study of language and behavior, or both, are in deep trouble, for reasons that should be familiar. The concept of language that Dummett takes to be essential involves complex and obscure sociopolitical, historical, cultural, and normative-teleological elements. Such elements may be of some interest for the sociology of identification within various social and political communities and the study of authority structure, but they plainly lie far beyond any useful inquiry into the nature of language or the psychology of users of language.

To take one example, consider the study of language acquisition. In ordinary usage, we say that a child of five and a foreign adult are on their way towards acquiring English, but we have no way to designate whatever it is that they "have." The child, in the normal course of events, will come to "have" English (at least partially and erroneously), though the foreigner probably will not. But if all adults were suddenly to die and children were somehow to survive, then whatever it is they are speaking would be a human language, though one that does not now exist. Ordinary usage provides no useful way to describe any of this, since it involves too many disparate and obscure concerns and interests, which is one reason why the concept of language that Dummett adopts is useless for actual inquiry. This matter is of some importance when we consider the reliance on notions of "misuse of language," "community norms," "social practice," and "rule following" that are often adopted as if they are sufficiently clear; they are not.[3]

In this connection, it is perhaps worthwhile to recall some further truisms; in rational inquiry, in the natural sciences or elsewhere, there is no such subject as "the study of everything." Thus it is no part of physics to determine exactly how a particular body moves under the influence of every particle or force in the universe, with possible human intervention, etc. This is not a topic. Rather, in rational inquiry we idealize to selected domains in such a way (we hope) as to permit us to discover crucial features of the world. Data and observations, in the sciences, have an instrumental character. They are of no particular interest in themselves, but only insofar as they constitute evidence that permits one to determine fundamental features of the real world, within a course of inquiry that is invariably undertaken under sharp idealizations,

often implicit and simply common understanding, but always present. The study of "language" in Dummett's sense verges on "the study of everything," and is therefore not a useful topic of inquiry, though one might hope, perhaps, to build up to a study of aspects of such questions in terms of what comes to be understood about particular components of this hopeless amalgam.

The conception of language as a "social practice" that Dummett and others propose raises further questions, as becomes clear when it is applied to concrete examples. Consider again examples (1) and (2) on page 47. In example (1), *feed herself* is taken to be predicated of Mary, but in example (2) it is predicated of some (female) person distinct from Mary; thus from example (2) it follows that I wonder which female person Mary expects to feed that very person, but not that I wonder which person Mary expects to feed Mary herself. The example raises many pertinent questions, among them, how we know these facts. The answer seems to be that the initial state of the shared language faculty incorporates certain principles concerning referential dependence (Binding Theory); and when certain options left undetermined in the initial state are fixed by elementary experience, then we have no more choice as to how to interpret examples (1) and (2) than we have about whether to perceive something as a red triangle or as a person. Social custom appears to have nothing to do with the matter in such cases, though in all of them, early experience helps set certain details of the invariant, biologically-determined mechanisms of the mind/brain. The same seems to be true rather generally. Taken literally at least, the proposals of Dummett and others concerning "social practice" appear to be false, as a matter of empirical fact. At the very least, some argument would be required to show why they should be considered seriously.

If language is construed as a social practice in the manner of these discussions, then it is tempting to understand knowledge of language as the learned ability to engage in such practices, as Dummett suggests or – more generally – as an ability that can be exercised by speaking, understanding, reading, talking to oneself, etc.: "to know a language just is to have the ability to do these and similar things" (Kenny 1984: 138).[4] The temptation is reinforced by a common construal of knowledge more generally as a kind of ability. This view contrasts with the conception of a language as a generative procedure that assigns structural descriptions to linguistic expressions, knowledge of language being the internal representation of such a procedure in the brain (in the mind, as we may say when speaking about the brain at a certain level of abstraction). From this point of view, ability to use one's language

(to put one's knowledge to use) is sharply distinguished from having such knowledge. The latter conception has two primary virtues:

1. It seems to be the right way to approach the study of human knowledge – knowledge of language in particular – within the general framework of the natural sciences, and it has proven a highly productive approach.
2. It is very much in accord with normal pre-analytic usage, a secondary but not entirely insignificant matter.

In contrast, the approach in terms of practical ability has proven entirely unproductive and can be sustained only by understanding "ability" in a way that departs radically from ordinary usage.

To see why this is so, suppose that Jones, a speaker of some variety of what we call "English" in informal usage, improves his ability to speak his language by taking a public-speaking course, or loses this ability because of an injury or disease (then recovers that ability, say, with a drug). Note that a speaker of "Japanese", under the same circumstances, would recover *Japanese*, not English, with the same drug, and plainly recovery in such cases differs radically from acquisition; a child could not acquire English or Japanese without any evidence. In all such cases, something remains constant, some property K, while ability to speak, understand, etc. varies. In ordinary usage, we say that K is knowledge of language; thus Jones's knowledge remained constant while his ability to put his knowledge to use improved, declined, recovered, etc. The account in terms of internal representation of a generative procedure accords with informal usage in this case. Note further that other evidence (say, from autopsy, were enough known about the brain sciences) might lead us to conclude that Smith, who never recovered English, not having taken the drug, nevertheless retained his knowledge of English intact after having completely lost his ability to speak and understand. (For more extensive discussion of these matters, and of possible alternative accounts, see Chomsky 1980; 1986.)

If knowledge is ability, then the property K must be a kind of ability, though plainly not ability in the quite useful normal sense of the word, since ability varied while K remained constant. We must therefore contrive a new technical sense of the term "ability," call it *K-ability*. Then K-ability remained constant while ability varied.[5] K-ability is completely divorced from ability, and has the properties of the old concept of knowledge; it might as well be called "knowledge," doctrinal matters aside.

It is rather ironic that these moves should be presented as in the spirit of the later Wittgenstein, who constantly argued against the practice of

constructing artificial concepts, divorced from ordinary usage, in defense of certain philosophical doctrines. In fact, the Wittgensteinian construal of knowledge as a species of ability seems to be a paradigm example of the practice that Wittgenstein held to be a fundamental source of philosophical error.

Notice that similar considerations show that *knowing-how* – for example, knowing how to ride a bicycle – cannot be analyzed in terms of abilities, dispositions, etc.; rather, there appears to be an irreducible cognitive element. Notice finally that an account of knowledge in terms of ability, taken in anything like its normal sense, has proven utterly unproductive. One might try accounting for the simple examples (1) and (2) in terms of Jones's abilities, for example. No such endeavor has ever been undertaken, and a close look at the problems makes it reasonably clear why it would have no hope of success.

The paradoxical flavor of ideas in the range I have been sampling becomes clearer when we look more closely at some of the specific injunctions. Take again Rorty's observation, taken as obvious without discussion, that "all the linguist has to go on is his observation of the way in which linguistic is aligned with non-linguistic behavior in the course of the native's interaction with the environment" (Rorty 1986: 339), apart from the "regulative principle" that the native informant is generally speaking truly. This conception, he notes, is drawn from Quine and Davidson. Thus in Quine's familiar paradigm of "radical translation" (Quine 1960; 1987), "field linguists" observing Jones must support their hypotheses *entirely* in terms of observation of Jones's behavior (or that of members of the "Jungle community," taken to be homogeneous; if it is not homogeneous, none of the arguments will go through, and if it is homogeneous, we may dismiss the community in favor of Jones without loss for these purposes, as I will do). I should note that in referring to Quine, textual questions arise, since – in response to queries and criticism – he has given many different versions of his paradigm, and these are not consistent (see Chomsky 1975: 187f., 198ff.). However, it is the one just cited, which Davidson and Rorty adopt, that is necessary if we are to be able to draw from Quine's paradigm any of the conclusions that are held to be important.

Before proceeding, let us note again that these prescriptions are radically different from the actual practice of the "field linguist." They are also completely foreign to the standard methods of the natural sciences. In the philosophical literature, the issues are generally discussed with regard to the theory of meaning and, in particular, with regard to aspects of the theory of meaning about which little is known (not, say, in connection with such matters as referential dependence, about which

a good deal is understood). This is dubious practice, because it means that controls on speculation by empirical knowledge and theoretical understanding are very slight. But if the doctrine has any validity, it should hold with regard to all of our attributions of linguistic competence, and Quine, at least, has been clear that this is so. Thus he explicitly argues that the same considerations hold when his "field linguist" alleges that in the sentence "John contemplated the problem" there are two phrases: the noun phrase *John* and the verb phrase *contemplated the problem*, not, say, the two phrases *John contemplated* and *the problem* or *John contemp* and *lated the problem*. According to Quine, at least when he is keeping to the assumptions required for his well-known conclusions to follow, this attribution of some property (knowledge, or whatever we choose to call it) to the informant Jones must be based exclusively on evidence about *Jones's behavior*; in fact, evidence used in accord with highly restrictive canons that he outlines. The same would also be true in the study of sound structure, relations of anaphors and antecedents, or whatever.[6]

It is worth noting that no linguist, or empirical scientist generally, would ever agree to be bound by such strictures. A comparable assumption in biology would be that in testing hypotheses about embryological development of humans, we cannot consider evidence obtained from the study of E. coli, or fruit flies, or apes, or physics. To mention one crucial case, in actual practice, every linguist approaches the study of a particular language on the basis of assumptions drawn from the study of other languages. Thus any linguist operating by the norms of the sciences would readily use evidence derived from the study of Japanese to help ground assumptions about Jones's knowledge of English. The logic is straightforward, and quite correct. There is overwhelming empirical evidence that people are not genetically "tuned" to acquire one rather than another language; rather, the "initial state" of their language faculty may be assumed to be uniform to a very good approximation. Presented with an array of evidence, the child acquires a specific language, making use of the resources of the initial state that determine a substantial part of the knowledge (competence) acquired; the initial state can be regarded as a fixed biologically-determined function that maps evidence available into acquired knowledge, uniformly for all languages.[7] Study of Japanese may, of course, provide us with evidence, perhaps compelling evidence, about the initial state, namely, by means of a comparison between what comes to be known and what is presented, the two being mediated by the resources of the initial state. If speakers of Japanese employ some formal property of language structure (say, *c-command*) in interpreting referential dependence, and the evidence available to the

Japanese child does not somehow "compel" or is not even conducive to this uniform result, we are entitled to attribute to the initial state a version of Binding Theory, incorporating this property and relevant principles involving it, and thus to explain the facts observed. But the initial state is shared by the English speaker Jones, and hypotheses about his initial state will of course have consequences as to the proper description of the cognitive state he attains. The conclusions derived from Japanese concerning Jones's knowledge of English might be far-reaching. Thus evidence about referential dependence in Japanese might prove relevant for determining the position of phrase boundaries in English.[8]

All of this is just standard scientific practice, never questioned – or even discussed, because it is so uncontroversial – in the natural sciences. However, Quine and those influenced by his paradigm are enjoining the "field linguist" to depart radically from the procedures of the sciences, limiting themselves to a small part of the relevant evidence, selected in accordance with behaviorist dogma; and also to reject the standard procedures used in theory construction in the sciences. The point is not academic: the normal practice of descriptive linguists crucially exploits these assumptions, which again should be the merest truisms.

We may put the point differently. The linguist and the child face radically different tasks. The child, endowed with certain innate capacities, acquires knowledge of a language – automatically, and with little if any choice in the matter. The linguist is trying to find out what knowledge the child acquires, and what innate properties of the mind/brain are responsible for this process of growth of knowledge (trying to find out what the child knows in advance of experience, to use a locution that seems to be quite appropriate). The linguist will quite properly use conclusions about innate properties, however derived, for the description of the knowledge attained, in particular, for the study of meaning, this domain having the same status as any other.

In fact, Quine's injunctions, consistently applied, would be still more extreme than this example indicates. Thus evidence from language pathology, or genetic variation, or neural structure, or biochemistry, or in fact evidence from any source, would be regarded by any scientist as potentially relevant in principle to determining the nature of the initial state or the state of knowledge attained, since these are simply elements of the natural biological world. Quine too insists on this point with regard to study of the natural world, apart from the study of humans above the neck when undertaken by "linguists," in his sense of this term. If it could be shown that some facts about the neural structure of the brain provide a natural realization of rule systems of one kind (say,

with the breakdown of "John contemplated the problem" into the two phrases *John* and *contemplated the problem*), but not other kinds, then this line of argument would be acceptable in the sciences to help settle the question of what is the correct description of Jones's knowledge – the cognitive state attained by Jones (the question of the choice of constituent structure in the case in question). The same is true with regard to the theory of meaning, or any empirical inquiry. But all of these paths, familiar in the natural sciences, are excluded by fiat under the Quinean conditions on the work of the "linguist" in accord with the paradigm that is widely adopted in the philosophical literature.

Quine has qualified these doctrines in interesting ways. A closer look at these qualifications reveals more clearly the arbitrary character of the stipulations imposed and the persistent misunderstanding of the empirical issues. As an example of arbitrary stipulation, consider Quine's discussion of the evidence that might lead us to assign one or another constituent structure to the sentences of Jones's English (Quine 1986). If this evidence derives from psycholinguistic experiments on perceived displacement of clicks,[9] then it counts; if the evidence derives from conditions on referential dependence in Japanese or on the formation of causative constructions in numerous languages, then it does not count – though this is evidence interpreted in the normal manner of the natural sciences, along the lines discussed a moment ago. Perhaps Quine might be interpreted as holding that evidence of the former type (so-called "psychological evidence") is in fact more powerful and persuasive than the so-called "linguistic evidence"; if so, this would simply be another error, since the opposite is the case, for the present at least. In fact, Quine appears to hold that the evidence differs in its epistemological character, a notion that is completely untenable. Evidence does not come labelled "for confirming theories" ("psychological evidence") or "for purposes of 'simplicity and general translatability'" ("linguistic evidence"). It is just evidence, good or bad, compelling or noncompelling, given the theoretical frameworks in which it can be interpreted for the purposes of sharpening or confirming hypotheses.

As an example of misunderstanding of empirical issues, consider Quine's discussion of the so-called "coordinate structure constraint," a descriptive generalization that covers, for example, the radical difference in status between the interrogative expressions derived by questioning "Mary" in the sentences "John saw Bill and Mary" and "John saw Bill with Mary": that is, the difference between "who did John see Bill and?," "who did John see Bill with?" Quine concludes that the "striking uniformity" exhibited in this constraint is not "a hint of a trait of all language," but "a hint of genetic kinship of the languages that seem

most readily grammatized in these terms."[10] This conclusion, however, is based on a serious misunderstanding of the empirical issues at stake. The problem is to explain how each child knows the relevant difference between "who did John see Bill and?" and "who did John see Bill with?" It cannot be that the child relies on evidence from the history of language, and the child typically has no relevant experience to determine (by "induction," or whatever) that the simple rule "Front-*wh*-phrase" is somehow blocked in the expression "John saw Bill and who" but not in "John saw Bill with who" (in colloquial English). Children do not, for example, produce "who did John see Bill and?," then to be informed by their parents that this is not the way it is done; and languages have not "drifted" to incorporate this "simplification" of the rule of question-formation over many millennia.[11] The problem, in short, is one of poverty of stimulus, and speculations about genetic kinship of languages have nothing whatsoever to do with it, in this and innumerable other similar cases.[12]

A similar refusal to permit the study of language to be pursued in the manner of the natural sciences is illustrated in other connections. Consider Davidson's article "A Nice Derangement of Epitaphs" in the volume cited earlier (Lepore 1986). Davidson considers the thesis that the goal of the descriptive study of meaning is to construct "an explicit theory" that "is a model of the interpreter's linguistic competence," a "recursive theory of a certain sort," and that we can "describe what an interpreter can do" only by appeal to such a theory. He then proceeds: "It does not add anything to this thesis to say that if the theory does correctly describe the competence of an interpreter, some mechanisms in the interpreter must correspond to the theory" (Davidson 1986b: 438). Similar points have been made by Dummett and others.[13]

For anyone approaching these problems from the standpoint of the natural sciences, the final comment quoted is utterly wrongheaded. If it had any validity, the analogous comment would apply in the study of visual perception, or chemistry. As elsewhere, it adds a great deal to the thesis to say that "some mechanisms in the interpreter . . . correspond to the theory." That is, natural scientists who construct a theory that "describes what an interpreter can do" will proceed to attribute to the subject certain fixed and explicit mechanisms that would have the properties assumed in this descriptive account, not others. The attribution might be at an abstract level, in terms of mentally-represented rule systems, or in terms of other abstract entities such as neural nets, or in terms of cellular structure, or whatever; all of this is standard natural science. Having proceeded to attribute specific structure and mechanisms to the person's mind/brain – often at some remove from unknown

"more elementary" physical mechanisms – the natural scientist is then in a position to test the theory in terms of a wide array of evidence, for example, evidence drawn from other languages in the manner just illustrated, or evidence from pathology or the brain sciences or biochemistry. Davidson's injunction blocks these efforts to employ the methods of rational inquiry in the sciences to determine whether the postulated account of the interpreter is indeed true, and to modify it if (as is likely) it is not.

The same problem arises when Quine, David Lewis (1983), Dummett, and many others object that some philosophical problem arises when linguists attribute to a speaker–hearer a specific internalized rule-system, and then seek to determine whether this theory of the person is true by the standard methods of the sciences. Perhaps this is even pure "folly," as Quine has argued (1972: 447), to be overcome by proper reflection on methodology. The perceived problem is that for a fixed array of observed behavior, or a fixed infinite set of utterances selected on some obscure basis and taken by the philosopher to be "the language," it is of course possible to construct infinitely many different theories that are consistent with this evidence ("grammars," as they are sometimes called); it is therefore held to be an unwarranted move to postulate that one of them is "true" and others "false" – unless, Quine sometimes holds, there is "psychological evidence" – with its mysterious properties that "linguistic evidence" lacks – to support one or another hypothesis. The argument is often buttressed by an analogy to the study of formal languages, which are completely irrelevant and highly misleading in this connection. If valid, the argument would hold throughout the sciences; in fact, it is nothing more than a form of skepticism that no one takes seriously in the study of the natural world for reasons that were clear by the seventeenth century, as Richard Popkin observes (Popkin 1979).[14] The natural scientist will attribute to the subject a specific system, not some other one (a "grammar," to use a misleading term), and will then proceed to determine whether this assumption is correct by seeking evidence of as wide a variety as possible, including crucially evidence from other languages, along the lines just discussed. Of course, there will always remain empirical indeterminacy, since this is empirical science, not mathematics, but that is all there is to say about the matter. A considerable literature exists arguing the contrary, but it is based on fundamental fallacies of reasoning.[15] Among these fallacies are the mistaken assumptions just discussed: that evidence about Jones's competence can only be drawn from Jones's behavior (interpreted in terms of the regulative principle about truth), and that it adds nothing to a description of Jones's behavior to attribute to Jones a

specific internal mechanism, perhaps a particular system of rules or some form of neural organization that realizes them.

The point can be illustrated, again, with the matter of phrase-structure boundaries. Suppose we have two kinds of evidence for the placement of the major boundary after the subject in "John – contemplated the problem," evidence from referential dependence in Japanese ("linguistic evidence") and evidence from perceptual displacement of clicks ("psychological evidence"). The first kind of evidence is subject to the familiar sort of indeterminacy. So is the second. Suppose that under experimental conditions established to yield the right results (typically, after many attempts that go wrong), clicks will be perceptually displaced to the subject–predicate boundary, not the verb–object boundary. These results can be interpreted as supporting the conclusion that the structure is [NP – V NP], not [NP V – NP] or [NP – V – NP]. But it is easy to apply Quine's argument to show that there is "no fact of the matter" in this case (Quine 1960: 303; see Chomsky 1980: 15). Plainly, there are many other interpretations of the experimental results. Perhaps clicks are perceptually displaced to the middle of a constituent, not its boundary; or perhaps the subject is responding by identifying the phrase-structure boundary directly below the major one. All other relevant experiments could be reinterpreted along similar lines, as can certainly be done in principle – though it is not so simple in practice, in the case of the "psychological" or "linguistic" evidence. The issues are the same throughout; or rather, there are no issues relevant here, since they hold of empirical inquiry generally.

When conclusions are drawn about phrase boundaries or other aspects of language on the basis of "linguistic evidence," Quine is reluctant to accept them "without further light on the nature of the supposed equipment,"[16] but when the same conclusions are based on "psychological evidence," these qualms do not arise. This epistemological dualism makes no sense whatsoever; it is a long step backwards from traditional metaphysical dualism, which was a rational reaction, on assumptions now known to be faulty,[17] to perceived empirical problems. The qualms, such as they are, are in principle the same, whatever the evidence on which conclusions are based, and are simply features of empirical inquiry. As for the "supposed equipment," it raises no problems of principle that differ from those characteristic of all theory construction in the empirical sciences.

Yet another paradox arises within this framework. Linguists, it is argued, are not permitted to attribute one particular language system rather than others to the individual or idealized community that they are studying;[18] they are not permitted to explore what is true of the brain, described at the level at which we construct rule systems and the like. But something is true of the brain; there is something about my brain that is more or less like yours and crucially different from the brain of a speaker of Swahili. Therefore someone should be permitted to study these aspects of the real world, but not linguists, who are restricted to inquiry into Jones's behavior and may not proceed to attribute specific mechanisms to Jones's mind/brain and to use evidence from other languages (or from any domain, in principle) to verify the accuracy of their conclusions about these mechanisms. Accepting these terminological strictures about what the linguist must do, the rational step is to abandon linguistics (including the study of meaning in accord with the conditions stipulated in the Quinean paradigm). Having abandoned these pointless pursuits, we may now turn to this other subject, where we are permitted to attribute specific mechanisms to Jones's mind/brain and to investigate these hypotheses by the methods of the sciences, using whatever evidence is at hand: in fact, the actual practice of linguists that is condemned in this curious, though extremely influential tradition in modern philosophy, which, in a final irony, prides itself on its "naturalism" and adherence to the methods of the sciences.

In his most recent efforts to justify the strictures he imposes, Quine (1987) offers the following argument. For the linguist, he argues, "the behaviorist approach is mandatory." The reason is that in acquiring language, "we depend strictly on overt behavior in observable situations . . . There is nothing in linguistic meaning, then, beyond what is to be gleaned from overt behavior in observable circumstances" (Quine 1987: 5), and the same holds true, by parity of argument, for the study of pronunciation, phrase structure, or whatever aspect of language we choose. Furthermore, as he makes explicit once again, the relevant behavior for the linguist is that of the natives to whom he or she is imputing knowledge of language: "If translators disagree on the translation of a Jungle sentence but no behavior on the part of the Jungle people [tacitly assumed to be homogeneous] could bear on the disagreement, then there is simply no fact of the matter," (Quine 1990: 38) and the linguist who holds that there are facts to be discovered, and that some theories (grammars) are correct and others not, is guilty of serious methodological error or pure "folly" (recall that the "translator" stands for the language learner as well[19] and that the same argument holds for pronunciation, phrase structure, etc.).

Consider now the following analogous argument. In reaching its final physical structure in the passage from embryo to mature state, the organism depends strictly on nutrition provided from outside (including oxygen, etc.). There is nothing in the physical structure of the mature

organism, then, beyond what is to be gleaned from the nutritional inputs. The student of human development and its outcome, then, must limit attention to these inputs; for the biologist, "the nutritionist approach is mandatory." The argument is the same as Quine's, and we see at once why it is untenable. True, the embryo "depends" on the nutritional environment just as the language learner "depends" on overt behavior. But what does the term "depends" include? Here we turn to the structure of the organism, which we may think of abstractly as a mapping M of external inputs into mature state. In the absence of such structure, observed behavior will lead to no knowledge of language and nutrition will lead to no growth. Quine of course recognizes this. Thus Quine's field linguist, pursuing the path of the language learner, "tentatively associates a native's utterance with the observed concurrent situation," and is permitted to make use of other hypotheses that allegedly correspond to capacities with which the language learner is endowed. If clarified, these hypotheses would constitute a theory of the innate structure of the organism and the mapping M.

As is agreed on all sides, without innate structure there is no effect of the external environment in language (or other) growth; in particular, without innate structure Jones could not have developed in a specific way from embryo to person, and his language faculty could not have assumed the state of mature competence that underlies and accounts for Jones's behavior. The child is endowed with this innate structure and therefore grows to maturity along a course that is largely inner-directed; the task of the scientist is to discover what the innate endowment is and what is the nature of the state attained. Currently, the best theory is that the initial state of the language faculty incorporates certain general principles of language structure, including phonetic and semantic principles, and that the mature state of competence is a generative procedure that assigns structural descriptions to expressions and interacts with the motor and perceptual system and other cognitive systems of the mind/brain to yield semantic and phonetic interpretations of utterances. A vast range of empirical evidence is relevant in principle to determining just how this proposal should be spelled out in detail. Again, all of this is normal science, yielding theories that are true or false[20] regarding Jones's competence and his initial state, part of the human biological endowment. Perhaps this approach should be abandoned in terms of some other conception, now unavailable; however, to establish this conclusion it does not suffice to demand that the linguist abandon the methods of the sciences.

As in his earlier formulations of these ideas, Quine's specific stipulations about the innate structure (hence the mapping M) are completely

arbitrary and, apart from their historical antecedents, here irrelevant. There is no reason to accept them in the case of language, just as comparable dogmatism about "dependence" would be rejected out of hand in the study of other aspects of the growth of organisms. Furthermore, there is compelling evidence that they are false, insofar as they are explicit. As in the study of physical development generally, the rational investigator will dismiss these dogmatic assumptions about the nature of "dependence" (that is about innate structure) along with other doctrines such as those just sketched, and will use whatever evidence can be found concerning the structure of the organism, the mapping M, and the nature of the states attained in particular cases. The conclusions that Quine, Davidson, Rorty and many others draw remain unargued. Nothing can be resurrected from the Quinean picture with regard to these matters, so far as I can see, though some of his conclusions – in particular, with regard to "meaning holism" – may well turn out to be correct, at least in large part.

Let us return now to the "analytic-synthetic" distinction, and the Davidsonian argument (Davidson 1986a: 313) that by "getting rid of it," Quine "saved philosophy of language as a serious subject." Recall that what is at issue here is not simply this distinction, but the question of language-determined semantic connections generally. As I mentioned, we cannot appeal to Rorty's argument, attributed to Quine, that the "field linguist" finds the distinction "of no use." In practice, semantic structure is regularly attributed to lexical items in descriptive work and theoretical studies on the semantics of natural language, and from these and other structural properties, semantic connections of various kinds are derivable, including analytic connections. There are good reasons for these standard assumptions about lexical structure. Acquisition of lexical items poses what is sometimes called "Plato's problem" in a very sharp form. As anyone who has tried to construct a dictionary or to work in descriptive semantics is aware, it is a very difficult matter to describe the meaning of a word, and such meanings have great intricacy and involve the most remarkable assumptions, even in the case of very simple concepts, such as what counts as a nameable thing. At peak periods of language acquisition, children are acquiring ("learning") many words a day, perhaps a dozen or more, meaning that they are acquiring words on very few exposures, even just one. This would appear to indicate that the concepts are already available, with much or all of their intricacy and structure predetermined, and that the child's task is to assign labels to concepts, as might be done with limited evidence given sufficiently rich innate structure. And these conceptual structures appear to yield semantic connections of a kind that will, in

particular, induce an analytic-synthetic distinction, as a matter of empirical fact.

To the extent that anything is understood about lexical items and their nature, it seems that they are based on conceptual structures of a specific and closely integrated type. It has been argued plausibly that concepts of a locational nature – including goal and source of action, object moved, etc. – enter widely into lexical structure, often in quite abstract ways. In addition, notions like actor, recipient of action, instrument, event, intention, causation and others are pervasive elements of lexical structure, with their specific properties and interrelations. Consider, say, the words *chase* or *persuade*. They clearly involve a reference to human intention. To chase Jones is not only to follow him, but to follow him with the intent of staying on his path, perhaps to catch him. To persuade Smith to do something is to cause him to decide or intend to do it; if he never decides or intends to do it, we have not succeeded in persuading him. Furthermore, he must decide or intend by his own volition, not under duress; if we say that the police persuaded Smith to confess by torture, we are using the term ironically. Since these facts are known essentially without evidence, it must be that the child approaches language with an intuitive understanding of concepts involving intending, causation, goal of action, event, and so on; furthermore, it must be that the child places the words that are heard in a nexus that is permitted by the principles of universal grammar, which provide the framework for thought and language, and are common to human languages as systems that enter into various aspects of human life. These elements also appear to enter into an integrated "conceptual scheme," a component of the initial state of the language faculty that is fleshed out in specific ways, with predetermined scope and limits, in the course of language growth, one aspect of cognitive development. There may be revision and restructuring of such conceptual schemes, (see Carey 1985), but care must be taken to separate out the various factors that enter into the course of development, including, quite possibly, genetically-determined maturation that yields effects perceived only in late stages of cognitive growth.

Notice again that we appear to have connections of meaning in such cases as these; we have a rather clear distinction between truths of meaning and truths of fact. Thus, if John persuaded Bill to go to college, then Bill at some point decided or intended to go to college and did so without duress; otherwise, John did not persuade Bill to go to college. Similarly if John killed Bill, then Bill is dead (though John may or may not be, depending on the facts). These are truths of meaning, not of fact. The *a priori* framework of human thought, within which

language is acquired, provides necessary connections among concepts, reflected in connections of meaning among words and, more broadly, among expressions involving these words, as in the example of referential dependence mentioned earlier. Syntactic relations provide a rich array of further examples. For example, there seems to be a clear distinction between the sentence "everyone who lives upstairs lives upstairs" and "everyone who lives upstairs is happy." Quine appears to believe that this distinction is more problematic and obscure than his distinction between "grammatical" and "ungrammatical," which he regards as somehow crucial for the linguist's investigations.[21] The opposite is the case. In fact, an absolute distinction between "grammatical" and "ungrammatical" appears to have little if any significance. It can be established one way or another or, perhaps better, not at all, since it is doubtful that the concept, in Quine's sense, plays any role in the theory of language. The reasons were discussed in the earliest work in generative grammar; this work is, in fact, the only work in which an effort was made to develop such a concept in some manner that might be relevant to linguistic theory, but in terms that were long ago understood to be inappropriate.[22]

It appears, then, that one of the central conclusions of modern philosophy is rather dubious: namely, the contention – often held to have been established by work of Quine and others – that one can make no principled distinction between questions of fact and questions of meaning, that it is a matter of more or less deeply held belief. This conclusion has been supported by reflection on an artificially narrow class of examples; among them concepts that have little or no relational structure. In the case of such sentences as "cats are animals," for example, it is not easy to find evidence to decide whether the sentence is true as a matter of meaning or fact, or whether there is an answer to the question in this case, and there has been much inconclusive controversy about the matter. When we turn to concepts with an inherent relational structure such as *persuade* or *chase*, or to more complex syntactic constructions such as those exhibiting referential dependence or causative and relative constructions, then it seems that semantic connections are readily discerned. Contrary to what Rorty and others assert, this is the common assumption of empirical work in the study of linguistic meaning, and, furthermore, it seems to be a reasonable assumption.

The status of a statement as a truth of meaning or of empirical fact can only be established by empirical inquiry, and considerations of many sorts may well be relevant; for example, inquiry into language acquisition and variation among languages. The question of the existence of analytic truths and semantic connections more generally is an

empirical one, to be settled by inquiry that goes well beyond the range of evidence ordinarily brought to bear in the literature on these topics. Suppose that two people differ in their intuitive judgments as to whether I can persuade John to go to college without his deciding or intending to do so (see Harman 1980). We are by no means at an impasse. Rather, we can construct conflicting theories and proceed to test them. One who holds that the connection between *persuade* and *decide* or *intend* is conceptual will proceed to elaborate the structure of the concepts, their primitive elements, the principles by which they are integrated and related to other cognitive systems, and so on; and will seek to show that other properties of language and other aspects of the acquisition and use of language can be explained in terms of the very same assumptions about the innate structure of the language faculty, in the same language and others, and that the same concepts play a role in other aspects of thought and understanding. One who holds that the connection is one of deeply held belief, not connection of meaning, has the task of developing a general theory of belief fixation that will yield the right conclusions in these and numerous other cases. Suppose one holds, with Paul Churchland for example, that the connection is based on the "semantic importance" of sentences relating *persuade* and *decide* or *intend* (that is, that these sentences play a prominent role in inference, or serve to introduce the term *persuade* to the child's vocabulary, and thus are more important than others for communication (Paul Churchland 1979: 51f.)). One then faces the task of showing that these empirical claims are in fact true. The first tack – in terms of innate conceptual structure – seems far more promising to me, and is the only approach that has any results or even proposals to its credit; it is, however, a matter of empirical inquiry, not pronouncements on the basis of virtually no evidence. Specifically, arguments against the first (conceptual) approach in terms of indeterminacy, unclarity, open issues, etc. establish nothing unless it is shown that alternative approaches in terms of some (now unavailable) theories of belief fixation or semantic importance are not subject to these problems.

The whole matter requires extensive rethinking, and much of what has been generally assumed for the past several decades about these questions appears to be dubious at best. There is, it seems rather clear, a rich conceptual structure determined by the initial state of the language faculty (perhaps drawing from the resources of other genetically-determined faculties of mind), waiting to be awakened by experience. All of this is much in accord with traditional rationalist conceptions and even, in some respects, the so-called "empiricist" thought of James Harris, David Hume, and others.

Many have found such conclusions completely unacceptable, even absurd; the idea that there is something like an array of innate concepts and that these are to a large degree merely "labeled" in language acquisition – as the empirical evidence suggests – certainly departs radically from many common assumptions. Some, for example Hilary Putnam, have argued that it is entirely implausible to suppose that we have "an innate stock of notions" including *carburetor* and *bureaucrat* (Putnam 1988a: 15). If he were correct about this, it would not be particularly to the point, since the problem arises in a most serious way in connection with simple words such as *table*, *person*, *chase*, *persuade*, *kill*, etc. However, his argument for the examples that he cites is not compelling. It is that to have given us this innate stock of notions, "evolution would have had to be able to anticipate all the contingencies of future physical and cultural environments. Obviously it didn't and couldn't do this" (p. 15).

Notice that the argument is invalid from the start. To suppose that, in the course of evolution, humans come to have an innate stock of notions including *carburetor* and *bureaucrat* does not entail that evolution was able to anticipate *every* future physical and cultural contingency – only these contingencies. That aside, notice that a very similar argument had long been accepted in immunology: namely, the number of antigens is so immense, including even artifically synthesized substances that had never existed in the world, that it was considered absurd to suppose that evolution had provided "an innate stock of antibodies"; rather, formation of antibodies must be a kind of "learning process" in which the antigens played an "instructive role." But this assumption might well be false. Niels Kaj Jerne won the Nobel Prize for his work challenging this idea, and upholding his own conception that an animal "cannot be stimulated to make specific antibodies, unless it has already made antibodies of this specificity before the antigen arrives" (Jerne 1985: 1059), so that antibody formation is a selective process in which the antigen plays a selective and amplifying role.[23] Whether or not Jerne is correct, he certainly could be, and the same could be true in the case of word meanings, the argument being quite analogous.

Furthermore, there is good reason to suppose that the argument is at least in substantial measure correct even for such words as *carburetor* and *bureaucrat*, which, in fact, pose the familiar problem of poverty of stimulus if we attend carefully to the enormous gap between what we know and the evidence on the basis of which we know it. The same is often true of technical terms of science and mathematics, and it surely appears to be the case for the terms of ordinary discourse. However surprising the conclusion may be that nature has provided us with an

innate stock of concepts, and that the child's task is to discover their labels, the empirical facts appear to leave open few other possibilities. Other possibilities (say, in terms of "generalized learning mechanisms") have yet to be coherently formulated, and if some day they are, it may well be that the apparent issue will dissolve.

In fact, it is not clear what thesis is being proposed by Putnam and others who reject what they call "the innateness hypothesis"; I should add that though I am alleged to be one of the exponents of this hypothesis, perhaps even the arch-criminal, I have never defended it and have no idea what it is supposed to be. Whatever the truth may be about antibody formation, it is based on the innate resources of the body and its immune system, and the task of the scientist is to find out what these resources are. Exactly the same is true of concept formation and language acquisition. For this reason, people who are supposed to be defenders of "the innateness hypothesis" do not defend the hypothesis or even use the phrase, because there is no such general hypothesis; rather, only specific hypotheses about the innate resources of the mind, in particular, its language faculty. General arguments against some unformulated "innateness hypothesis" have no bearing on actual hypotheses about innateness, in the case of growth of language and conceptual systems or other forms of physical growth.

Putnam offers a counter-argument to the one just sketched on analogy to the immune system. He points out that concepts "often arise from *theories*," and the number of possible theories (or perhaps even "theory *types*") is so immense, even for "short" theories, as to make "the idea that evolution exhausted all the possibilities in advance wildly implausible" (Putnam 1988a: 128). The argument is correct, but again irrelevant. In the first place, we are considering what humans are capable of acquiring, and there is no reason to believe that "all theories" can be learned or constructed by humans, nor is it even clear what sense this thesis has.[24] Furthermore, Putnam's original argument was supposed to bear on the specific words *carburetor* and *bureaucrat*, and no cardinality argument is relevant to these cases, or to any substantive empirical hypothesis about innate structure. In other words, his argument that "evolution couldn't have done that" simply does not hold in the cases for which it is offered. The argument that evolution couldn't have done "everything" – even what is beyond human capacity – might hold if one could make some sense of it; such an argument would not, however, be relevant here, even if it could be given in a coherent form.

In the same connection, Putnam argues that the thesis of "meaning holism," with the Quinean principle that "revision can strike anywhere," contributes to undermining certain conclusions concerning the innate

structure of conceptual systems and language generally. But this line of argument is questionable. Suppose that the thesis of "meaning holism" is correct in the sense that, as Putnam puts it, there are no "'psychologically real' entities which have enough of the properties we preanalytically assign to 'meanings' to warrant an identification," and reference is fully determined only on holistic grounds. Nevertheless, it does not follow that semantic connections cannot be completely fixed and stable as a matter of biological endowment. Thus certain relations may remain stable as other considerations lead to various choices about fixing of reference. Furthermore, empirical considerations of the kind discussed earlier bear on the question of whether it is indeed true that "revision can strike anywhere." The point cannot be established for natural language by reference to the practice of the natural sciences from which Putnam draws many of his examples; these arguments, assuming them to be correct, do not suffice to show the absence of intrinsic semantic and conceptual structure based on fixed properties of the human mind. The thesis of "holism" may be correct in some measure or form, but the questions of semantic connections in natural language remain to be settled by empirical study, and – for the present at least – the evidence appears to support their existence – rather strongly, it seems to me.

Let us pursue further Davidson's argument in his paper "A Nice Derangement of Epitaphs," (1986b) in which he purports to show that the study of actual communication undermines a "commonly accepted account of linguistic competence and communication" and shows that "there is no such thing as a language, not if a language is anything like what many philosophers and linguists have supposed. There is therefore no such thing to be learned, mastered, or born with" (Davidson 1986b: 446). This conception of language, which Davidson believes to be refuted, is founded on three basic assumptions concerning what he calls "first language" or "prior theory," a "complex system or theory" shared more or less by speaker and hearer (p. 436). The assumptions are:

1. that the prior theory is "systematic" in the sense that the *interpreter* who has this theory is able to interpret utterances on the basis of properties of their parts and the structure of the utterance;
2. that this method of interpretation is shared; and
3. that the component elements of the system are governed by learned conventions or regularities.

The third of these assumptions is untenable for other reasons, but instead of delaying on this matter, let us present it in the form required

for Davidson's argument: the component elements of the system are available, as he puts it, "in advance of occasions of interpretation"; it is a fixed element in communication situations, for interpreters at a fixed state of language knowledge.

To refute this conception, Davidson observes that in ordinary communication situations the interpreter makes use of all sorts of conjectures and assumptions about what the speaker may have in mind, relying on properties of the situation, the speaker's presumed intentions, and so on. The interpreter thus "adjusts his theory," modifying the "prior theory" to a "passing theory" that is "geared to the occasion." But this "passing theory cannot in general correspond to an interpreter's linguistic competence." This "passing theory is not a theory of what anyone (except perhaps a philosopher) would call an actual natural language" (Davidson 1986b: 443), Davidson continues, and "'Mastery' of such a language would be useless, since knowing a passing theory is only knowing how to interpret a particular utterance on a particular occasion" (p. 443). Furthermore, communication can proceed quite well when the prior theory is not shared by speaker and hearer, and the prior theory too is not what "we would normally call a language" since it is a psychological particular, specific to the speaker-hearer with features that are not shared through the "community." The interpreter has some kind of "strategy," a "mysterious process by which a speaker or hearer uses what he [or she] knows in advance plus present data to produce a passing theory," and for communication, what two people need "is the ability to converge on passing theories from utterance to utterance." Given these facts, there is no longer any use for "the concept of a language," for "shared grammar or rules," for a "portable interpreting machine set to grind out the meaning of an arbitrary utterance"; rather, we need something more evanescent, mysterious and "holistic," "the ability to converge on a passing theory from time to time" (p. 445). We thus are led to "abandon . . . not only the ordinary notion of a language, but we have erased the boundary between knowing a language and knowing our way around in the world generally . . . In linguistic communication nothing corresponds to a linguistic competence" (pp. 445–6) based on the three principles just mentioned, because "there are no rules for arriving at passing theories." At the conclusion of the discussion, however, Davidson asserts that a passing theory is derived somehow "from a private vocabulary and grammar," that is, from a "prior theory" meeting the first and perhaps a version of the third condition, but possibly not shared in the "community"; there is then a "prior theory" and there are surely certain methods, not others, "for arriving at passing theories," whether or not one wants to call these methods "rules" (p. 446).

The various parts of the argument are largely correct, but they do not seem to show very much. In particular, no reason has been offered to doubt that there is a "prior theory" in the usual sense of the study of language and knowledge of language; that is, a specific generative procedure incorporated in a specific mature state of the language faculty. Of course, this "prior theory" will be quite different from what is called "a language" in ordinary usage, but this is because no such concept plays a role in empirical inquiry into language and mind, as already noted.

In the face of Davidson's arguments, we may continue to suppose that there is, to very good first approximation, a fixed and invariant language faculty which maps presented evidence onto a system of rules and principles (or whatever turns out to be correct with regard to the cognitive state attained) that assign interpretations to utterances. Call this acquired system a "generative procedure." To know a language is to have an internal representation of this generative procedure, which we will express at various levels of abstraction from "more elementary" mechanisms and will seek to relate to such mechanisms, in the normal manner of the natural sciences.[25] Proceeding in accord with normal practice, we may also seek to construct a "parser" – a device, also attributed to the mind/brain – which incorporates the generative procedure attained along with other specified structures and properties,[26] and maps presented utterances into structural descriptions that are interpreted by other components of mind. So far, we are dealing with feasible questions of empirical inquiry.

There is also a further problem, which we can formulate in vague terms but which cannot be studied in practice: namely, to construct an "interpreter" which includes the parser as a component along with all other capacities of the mind – whatever they may be – and accepts nonlinguistic as well as linguistic inputs. This interpreter, presented with an utterance and a situation, assigns some interpretation to what is being said by a person in this situation. The study of communication in the actual world of experience is the study of the interpreter, but this is not a topic for empirical inquiry, for the usual reasons: there is no such topic as the study of everything. Similarly, science does not investigate other phenomena of the world as presented to us in everyday experience. The interpreter – as Davidson correctly observes – includes everything that people are capable of doing, which is why it is not an object of empirical inquiry, and why nothing sensible can be said about it. We might hope to learn something about various elements of the interpreter, proceeding by the normal methods of the sciences, beginning with the "private vocabulary and grammar" that constitute the language attained, proceeding to the parser, then perhaps – to the extent feasible – turning

to other elements of the mind and of situations that enter into normal human life. However, if we begin with the demand for a theory of everything, we will find nothing; it is unnecessary to construct elaborate arguments to establish this point.[27] The situation is no different in the far more advanced sciences. The proper conclusion is not that we must abandon concepts of language that can be productively studied, but that the topic of successful communication in the actual world of experience is far too complex and obscure to merit attention in empirical inquiry, except as a guide to intuitions as we pursue research designed to lead to some understanding of the real world, communication included. These observations have no bearing on whether or not there is a "prior theory," that is, an internalized generative procedure, in the normal sense of empirical practice.

Davidson's "passing theory" is not a useful notion; about this, he is surely correct. The interpreter will construct all sorts of "passing theories" (though, crucially, not *any* sort), changing moment to moment, because the interpreter as Davidson conceives it includes everything available to human intelligence; it makes no sense, however, to call its transient states "theories" or to consider them a subject of direct inquiry. Crucially, nothing in Davidson's argument bears on the assumption that the "prior theory" (though not understood quite in his terms) remains a fixed and invariant element of the "interpreter" (as of the narrower idealized parser), and that it enters into the functioning of the interpreter.

In this discussion, Davidson focuses attention on malapropisms and so-called "misuse of language" more generally. Here some care is necessary. Let's again take Jones, a speaker of a variety of what we informally call "English." Jones has mastered a generative procedure that associates with utterances structural descriptions, including semantic properties, and has other capacities of mind that allow him to produce and interpret linguistic expressions making use of these structural descriptions. Let us call this generative procedure his "I-language," where *I* is to suggest "internalized" (in the mind/brain) and "intensional" (in that the procedure is a function enumerating structural descriptions, considered in intension with a particular description).[28] Here we are referring to specific postulated mechanisms of the mind/brain, considered abstractly.

Jones may speak in a way that is not in accord with his I-language, or may offer judgments inconsistent with his I-language; judgments about ourselves, like others, can be mistaken, and much more than I-language is involved in behavior. This is an uninteresting case of misuse of language; call it the "individual sense."

Suppose that Jones, like most of us, normally says such things as "hopefully, we'll be able to solve that problem," or uses the word

"disinterested" to mean uninterested. Various authority figures tell us that this is "wrong," a "mistake," or not in accord with the "rules of English." Jones is "misusing his language," namely English, a language of which he has only a partial and perhaps distorted knowledge, as in Dummett's "fundamental sense" of language. Even if 95 per cent of the population – or for that matter everyone but William Safire and a few others – were to behave in the manner of Jones, these cases would still constitute "misuse of language." Or Jones may try to adapt to the practice of some community for some reason, or perhaps for no reason at all, and may fail to do so, in which case people observing Jones may speak informally of a misuse of the language of this community. These concepts of "misuse of language," which we may call "the community sense," may be of interest for the study of the sociology of group identification, authority structure, and the like, but they have little bearing on the study of language, so far as we know. We understand this perfectly well in the case of pronunciation. Thus to say that one variety of English is "right" and another "wrong" makes as much sense as saying that Spanish is right and English wrong; and the same is true – though for some reason the point seems more obscure – with regard to other aspects of language.

Another possible sense of the concept "misuse of language" derives from Hilary Putnam's notion of "the division of linguistic labor." Thus in the lexicon represented in my mind/brain, the entry for "elm" and "beech," or "mass" and "kinetic energy," may include an indication that the reference for these terms is to be determined by experts to whom I defer. Then I might apply the terms inaccurately, in the sense that the reference is not in accord with the determinations of these experts. In this case, I might be said to be "misusing my own language."[29] Let us call this the "expert sense" of misuse of language. Again, nothing of great moment appears to follow, surely nothing relating to the approach to language within the framework of individual psychology sketched earlier, and typically followed in practice.[30] Notice that no useful concept of "language" or "community" emerges from these considerations. Thus my expert for "elm" and "beech" may be an Italian gardener who speaks not a word of English, and who corrects my usage through reference to the technical Latin names that we share; and my expert for "mass" and "kinetic energy" may be a monolingual German physicist. But we would not conclude that German and Italian are included in English, or that all of us form a "community" in any useful sense of the term.

Is there any other concept of "misuse of language"? I am aware of none. If so, the concept plays no important role in the study of language, meaning, communication, or whatever. To take some examples

of the kind that Tyler Burge has discussed, suppose that Jones uses the term "arthritis" to refer to a pain in the thigh. Suppose this is the usage of his village, but not the usage of the outside community. Jones is not misusing his language in the individual sense; his usage is true to his I-language. In his village, he is not misusing his language in the community sense, but outside its borders, he is. Depending on how "arthritis" is represented in Jones's mental lexicon, he may or may not be misusing his language in the "expert sense." How should we attribute beliefs about arthritis to Jones? Here intuitions differ, and it may be that evidence is too slim, for the moment, to settle the point satisfactorily. Putting aside the "expert sense," suppose we use the term "I-belief" to refer to the concept that is like belief, except that Jones has the same belief within his village and in the wider community, namely, the belief that we would express, in our I-language, by saying that he has some kind of body pain.[31] This may or may not be the same as the concept of belief in our ordinary language, but it is the concept that seems to be required for the study of what is misleadingly called "the causation of behavior" – misleadingly, because it is unclear that behavior is "caused" in any useful sense of the term. Clearly, there is no reason to suppose that the concepts of general psychology will be those of ordinary usage, just as the concepts of physics, or of the subbranch of psychology called "linguistics," typically are not. Nor is it at all obvious to me that there is a reasonable branch of science (or to be more accurate, human science, meaning the kind of scientific inquiry that humans, with their particular cognitive capacities, are capable of undertaking) that deals with questions of this nature.

It has not, I think, been established that there is anything more to say about the matter. In particular, reference to "misuse of language," to "norms," to "communities," and so on seems to me to require much more care than is often taken. These concepts are obscure, and it is not clear that they are of any use for inquiry into language and human behavior. Any argument that relies on these notions merits careful scrutiny, and I doubt that familiar arguments can withstand it. Communities are formed in all sorts of overlapping ways, and the study of communities and their norms quickly degenerates into the study of everything. The fact remains that Jones speaks and understands the way he does on the basis of the I-language he has acquired in the course of language growth; and if Jones does or does not follow what we choose, for some transient purpose, to call "community norms" or "social practice," it is on the basis of this internalized I-language (along with much else). Boris, a monolingual speaker of some variety of Russian, has a different I-language, and follows different "norms." I can understand

Jones, within limits, because my I-language is not too different from his, and because he and I more or less share other unknown properties that enter into the full interpreter; this is not a topic of empirical inquiry as it stands, in its unanalyzed complexity. That seems to me the way we should approach these questions.

In these terms, we can develop a concept of "knowledge of language" that is appropriate for the inquiry into language and mind; namely, mastery and internal representation of a specific I-language. The linguist's grammar is a theory of the I-language, and universal grammar is the theory of the initial state of the language faculty. Jones's I-language is one particular mature state – or output, regarding the language faculty as a function that maps evidence into I-language. What about the concept language? We might simply understand languages as I-languages, thus taking a language to be something like "a way of speaking," the "finite means" that provide for "infinite use" in the terms of Wilhelm von Humboldt's characterization of language (1836: 122, paragraph 13; 1988: 91; see also Chomsky 1964: 17), also an effort to capture his concept of language as a "process of generation" rather than a set of "generated objects." We thus take language to be, in effect, a "notion of structure" that guides the speaker in forming "free expressions," in Otto Jespersen's terms (1924: 19; see also Chomsky 1977). For empirical inquiry, I think that is an appropriate decision, though obviously not for ordinary discourse. Alternatively we might want to construct a concept of language divorced from cognitive states, perhaps along lines suggested by James Higginbotham (1989). Taking knowledge of language to be a cognitive state, we might construe the "language" as an abstract object, the "object of knowledge," an abstract system of rules and principles (or whatever turns out to be correct) that is an image of the generative procedure, the I-language, represented in the mind and ultimately in the brain in now-unknown "more elementary" mechanisms. Since the language in this sense is completely determined by the I-language, though abstracted from it, it is not entirely clear that this further step is motivated; perhaps it is, however.

In these terms, it seems to me that the questions about language and its use that can be subjected to empirical inquiry can readily be formulated, and as far as we now know, best addressed. There may well be many other questions that are not subject to empirical inquiry in the manner of the sciences – and perhaps never will be – if humans are themselves part of the natural world, and thus have specific biological capacities with their scope and limits, like every other organism. We must be careful not to succumb to illusions about evolution and its adaptive miracles. There is nothing in the theory of evolution that

suggests that we should be able to answer questions that we can pose, even in principle, even if they have answers, or that we should be able to pose the right questions. To the extent that we can, we have empirical science, a kind of chance convergence of properties of the mind and properties of the extra-mental world. There is nothing surprising about this; we take for granted that something similar is true of rats and bees, and should not be surprised to learn that humans are biological organisms, not angels. Within the limits of human science, however, it seems to me that the best guess as of the present is that the framework I have just briefly outlined is a proper one for inquiry into the empirical questions about language and mind; and within it, there are some notable successes and many intriguing prospects.

4 Naturalism and dualism in the study of language and mind

The terms of the title can be understood in various ways, along with the frameworks in which they are embedded. I would like to outline interpretations that I think are useful and proper, and to suggest a more general thesis, which would require much more comprehensive argument: that there is no coherent alternative to proceeding in this way for the range of issues addressed, and that other endeavors in roughly the same realm are clarified and facilitated if understood as extensions of the approach outlined.

Deflating the terms

Putting "language" aside for the moment, let's begin by taking the other terms of the title in ways that are innocent of far-reaching implications, specifically, divorced from any metaphysical connotations. Take the term "mind" or, as a preliminary, "mental." Consider how we use such terms as "chemical," "optical," or "electrical." Certain phenomena, events, processes, and states are called "chemical" (etc.), but no metaphysical divide is suggested by that usage. These are just various aspects of the world that we select as a focus of attention for the purposes of inquiry and exposition. I will understand the term "mental" in much the same way, with something like its traditional coverage, but without metaphysical import and with no suggestion that it would make any sense to try to identify the true criterion or mark of the mental. By "mind," I mean the mental aspects of the world, with no concern for defining the notion more closely and no expectation that we will find some interesting kind of unity or boundaries, any more than elsewhere; no one cares to sharpen the boundaries of "the chemical."

Furthermore, I keep here to the human mind (visual system, reasoning, language, etc.). There is no quest for a unified science of locomotion, ranging from amoeba to eagle to science-fiction spaceship; or of communication, ranging from cell to poetic discourse to imagined extraterrestrials. Rather, biologists study how dolphins swim and ants communicate,