INFORMATION TO USERS

This was produced from a copy of a document sent to us for microfilming. While the most advanced technological means to photograph and reproduce this document have been used, the quality is heavily dependent upon the quality of the material submitted.

The following explanation of techniques is provided to help you understand markings or notations which may appear on this reproduction.

1. The sign or "target" for pages apparently lacking from the document photographed is "Missing Page(s)". If it was possible to obtain the missing page(s) or section, they are spliced into the film along with adjacent pages. This may have necessitated cutting through an image and duplicating adjacent pages to assure you of complete continuity.

2. When an image on the film is obliterated with a round black mark it is an indication that the film inspector noticed either blurred copy because of movement during exposure, or duplicate copy. Unless we meant to delete copyrighted materials that should not have been filmed, you will find a good image of the page in the adjacent frame.

3. When a map, drawing or chart, etc., is part of the material being photographed the photographer has followed a definite method in "sectioning" the material. It is customary to begin filming at the upper left hand corner of a large sheet and to continue from left to right in equal sections with small overlaps. If necessary, sectioning is continued again—beginning below the first row and continuing on until complete.

4. For any illustrations that cannot be reproduced satisfactorily by xerography, photographic prints can be purchased at additional cost and tipped into your xerographic copy. Requests can be made to our Dissertations Customer Services Department.

5. Some pages in any document may have indistinct print. In all cases we have filmed the best available copy.

University Microfilms International

300 N. ZEEB ROAD, ANN ARBOR, MI 48106
18 BEDFORD ROW, LONDON WC1R 4EJ, ENGLAND

8014986

SAMET, JERRY HOWARD

TACIT KNOWLEDGE

City University of New York PH.D. 1980

University Microfilms International 300 N. Zeeb Road, Ann Arbor, MI 48106 18 Bedford Row, London WC1R 4EJ, England

Copyright 1980

by

Samet, Jerry Howard

All Rights Reserved

TACIT KNOWLEDGE

BY

JERRY SAMET

A dissertation submitted to the Graduate Faculty in Philosophy in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York.

1980

© COPYRIGHT BY

JERRY SAMET

1980

This manuscript has been read and accepted for the Graduate Faculty in Philosophy in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

Feb. 1, 1980
date

Harold Zvalow
Chairman of Examining Committee

2-1-80
date

Marshall Cohen
Executive Officer

Prof. Robert Fiengo

Prof. Jerrold Katz

Prof. Robert Schwartz
Supervisory Committee

The City University of New York

PREFACE

I originally became interested in the topic of this study by way of a problem which seems to me to be at the juncture of a number of intellectual disciplines. The problem is this. Some of us know English, some of us don't; how do we explain this discrepancy? This is admittedly not much of a puzzle, in that many straightforward answers come to mind immediately. It is usually enough to point out that what's special about the people who know English is that they've learned it. The ones who don't know English probably learned some other language. Knowing English is in this sense not like having perfect pitch. But then again, it's not like what you know when you've learned the names of the capitals of the fifty states either. English is not a closed body of specific facts that can be committed to memory and reproduced at will one at a time. So what did the ones who know English learn that the others didn't learn? In the light of this question we might try to go beyond the vague reference to learning and say that the ones who know English are just those who have learned and now know the grammatical structure of the language and the meanings of the words. Because they lack this crucial knowledge, the others don't know English.

I think we have a firm conviction that something like this has to be right. But do English speakers really know these things? Putting aside the question of specific word meaning for the time being, there is reason to wonder whether speakers really know anything about the structure of their language. After all, isn't this what linguists try to discover? If speakers had such knowledge, then they could just deliver it into the linguist's hands. This situation has led to the proposal that whether they realize it or not, speakers of a language have such knowledge of underlying principles of language. One can say that their knowledge is in this sense tacit: not only are they generally unaware of having such knowledge, but they cannot become aware of it by introspecting.

The claim that speakers have such knowledge combines two controversial views. From a general epistemological perspective, it involves the idea that there can be knowledge that is more or less permanently hidden from the awareness of the knower; we might say that such knowledge is 'buried'. Apart from its epistemological novelty, the tacit knowledge hypothesis also introduces a new treatment of linguistic functioning as involving unconscious rule application. In this dissertation I defend the notion of tacit knowledge by arguing that both these views are essentially correct. My main interest,

however, is in the general epistemological claim; I concentrate on its application to the language case because most of the relevant discussions of the issues have been made in relation to our knowledge of language. If it can be shown that we have some sort of underlying tacit knowledge of linguistic principles, then a fortiori we have tacit knowledge.

In Chapter I I discuss the locus modernus of the tacit knowledge controversy--viz. Noam Chomsky's Aspects of the Theory of Syntax. I then present a set of conditions adapted from Brentano's attack on unconscious ideas that I suggest should be satisfied by any attempt to show that tacit knowledge ascriptions are legitimate.

In Chapter II I consider a number of theoretical objections to the view that there can be any such thing as knowledge that is tacit. These objections have two different but related foundations. The first set is based on the view that tacit knowledge, as I have defined it here, fails to meet certain a priori criteria that anything which is bona fide knowledge must satisfy. These criteria are in some wise to be connected to the Freudean notion(s) of the unconscious. I argue that these objections can be met and are based on a misunderstanding of the Freudean position. The second set of

objections is based on the view that there are insurmountable methodological problems relating to the notion of extensional equivalence that prevent us from ascribing tacit knowledge. Here I argue that the problem is no worse in this case than in the ordinary ascription of knowledge and belief. The chapter concludes with a discussion of Quine's position on tacit knowledge ascriptions.

Chapter III is divided into two parts. In the first, I discuss a number of further arguments against the notion of tacit knowledge. Although in themselves unsuccessful, these arguments point the way to what I take to be a valid objection to the Chomskyan position--specifically, the problem of constraining ascriptions of tacit knowledge. The second part of the chapter takes up Jerry Fodor's conception of tacit knowledge as a response to this last objection. Fodor's view is that we should ascribe tacit knowledge to an organism on the basis of the internal workings of an optimal simulation of the organism. I argue that Fodor fails to provide a convincing argument for this proposal; the arguments that can be extracted from his papers lead to implausible results.

Chapter IV has three parts. In the first, I consider another model for tacit knowledge ascriptions proposed by Graves, Katz, et. al.. This conception I argue, meets the objections levelled against the Chomskyan position. The key to this proposal is that we only ascribe tacit knowledge in cases where we have some unexplained explicit knowledge to account for. In the next section I consider and reject a number of arguments against the Graves-Katz approach. Finally I argue that two proposed alternatives to explanation by ascription of tacit knowledge are unacceptable.

After I'd completed this work I noticed my tendency to discover at least two different possible positions (or arguments) hidden in every one that I discussed. There are two possible explanations: (1) that I have a tendency to create an 'alter-argument' to go along with every real one, and (2) that others have a tendency to run together different strands of thought headed in the same direction. If the former explanation is correct I recommend that the reader imagine that the alter- argument is a hypothetical one from some 'hypothetical objector'. If the second explanation is right, then I certainly have no need to apologize.

The writing of this thesis has stretched over too long a period, and, much to the chagrin of my wife and family, the bulk of it has been rewritten very recently. I hope that any disparity in style is not distracting.

My greatest indebtedness for this work (and for much else besides) is to my advisor Jerry Katz. He encouraged me from the start and guided me through many, many difficult periods. It will be obvious that many of my ideas and discussions bear his stamp (though perhaps not his approval) and it would only be clumsy to try and credit every suggestion he made. I have benefitted more from his friendship and instruction than I can say and I welcome this opportunity to express my gratitude.

Many of the ideas developed here arose out of discussions with my brother-in-law Alan Zaitchik who has been a constant source of inspiration for me. The second half of Chapter II owes much to Alan Berger's attempts to set me straight on Quine's views. He is (in my eyes) not morally responsible for any/all of the misunderstandings and misinterpretations that I still stubbornly cling to. The first two chapters were improved by Peter Lupu's careful comments and criticism.

I also want to thank the administration, faculty, and students of the CUNY Graduate Center for their friendship and support. Special thanks are due to Michael Levin for his consistently provocative opposition to almost all my positions and careful attention to all my ideas, to Stefan Baumrin for teaching me how to approach a philosophical text, and to Howard Platzman for an indefinite number of enlightening and enjoyable conversations.

I owe a great deal to my new family and old and dear friends the Zaitchiks--Samuel, Nell, Linda, and Toby Kaufman--who went through the worst with me and will unfortunately not have a thesis to show for it.

Finally, to my sister Sharon and her husband Yaakov, and to my parents Imre and Rose Samet whose confidence, patience, and love sustained me throughout, my deepest thanks.

This work is dedicated with all my love to my wife Deborah for her sharp critical mind, warm heart, and blind faith.

TABLE OF CONTENTS

IV. THE INFERENTIAL MODEL

CHAPTER I

THE CHOMSKYAN BACKGROUND

My two main purposes in this brief chapter are first, to say something about what tacit knowledge is and is not, and to show how the claim that we have such knowledge arises in Chomsky's conception of linguistics. Second, I want to map out my general line of defense of the notion of tacit knowledge and to indicate some of the problems that must be handled if the defense is to be successful. The first section pinpoints what I _don't_ have in mind in talking about tacit knowledge. The second section is a discussion of the relevant features of Chomsky's methodological suggestions in _Aspects of a Theory of Syntax_. In the final section I consider Brentano's proposed conditions of adequacy on ascriptions of unconscious mental states and adapt them to the present framework.

The Via Negativa: What Tacit Knowledge Is Not

Epistemologists have traditionally suffered from a strange sort of paranoia--they stand in constant fear of skepticism. For Descartes such a fear may have been well-founded and realistic; for most of us it is not.[1]

My concern in this thesis will be with the other side of the skeptical coin. The skeptic wants to claim that we do not know as much as we think we do; my own view is that he couldn't be more wrong. I will argue that we know more than we think we do; that in addition to the things we know that we are aware of knowing, we possess a body of tacit or implicit knowledge.

I want to make sure that my position is not misconstrued as being identical with either of two related views to which it bears some resemblance. First, someone might say that I am attacking a straw man. Surely, everyone grants that of the great many things which we have learned and subsequently forgotten, some items can be recovered when the need arises. We often surprise ourselves and others with what we know. We play a song at the piano quite fluently even though we thought we had long ago forgotten the right chords; we suddenly know that Kant died in 1804 though we have just denied having any detailed knowledge of pertinent dates in the history of philosophy ("I must have read that somewhere, I guess."); and so on. This is an innocent sense of the term 'tacit', and taken in this way it is uncontroversially true that some of our knowledge <u>is</u> tacit. The cases I have in mind go beyond this; they differ from the Kant and piano examples in at least one crucial way. In the latter sort of case we are usually aware of ac-

quiring the knowledge (at the time of acquisition) and the problem is one of memory. In the cases I want to defend, memory does not seem to be the culprit. These are cases where the subject is unaware of having ever acquired the knowledge in the first place.

The second related view is typified by the geometrical knowledge that Socrates tries to elicit from Meno's slave. Here too we have the idea that some of our knowledge is implicit or not conscious. Unlike the 'innocent' cases of the previous paragraph Plato does not assume that the knowledge in question was consciously acquired, or that the implicitness is attributable to a simple lapse of memory. Nevertheless, there is still an important difference between the Meno-type cases and the ascriptions of tacit knowledge I want to defend. For Plato, although the slave's knowledge of geometry is not conscious at the beginning of Socrates' questioning, it is at least potentially conscious. This, I take it, is part of the point of the dialogue: that Socrates' skillful elicitation will enable the slave to remember what was formerly implicit in him. I will claim that we have knowledge that fails to meet even this elicitation condition, so I will reserve the term 'tacit' for knowledge that is not even potentially conscious in this sense.

I now have to explain why anyone would want to defend so cantankerous a view. The position I will be adopting has its roots in Chomsky's conception of the relation between a speaker-hearer[2] of a language and the grammar of that language that the linguist tries to discover. As far as it is pertinent to our concerns, Chomsky makes the following two controversial claims: first, that a grammar is a psychological theory whose goal it is to characterize the rules of language that every speaker has internalized; and second, that our knowledge of a language is to be explained in terms of our tacit knowledge of these underlying rules. These claims are embedded in a radically new framework that recasts the relationship between linguistics, on the one hand, and psychology and epistemology, on the other.

The Methodological Revolution

Grammar as psychology

The first chapter of Aspects contains the following three claims:

> A grammar of a language purports to be a description of the ideal speaker-hearer's intrinsic competence.
>
> We thus make a fundamental distinction between competence (the speaker-hearer's knowledge of his language) and performance (the actual use of the language in concrete situations).

> The problem for the linguist, as well as for te child learning the language, is to determine from the data of performance the underlying system of rules that has been mastered by the speaker-hearer and that he puts to use in actual performance. Hence, in the technical sense, linguistic theory is mentalistic, since it is concerned with discovering a mental reality underlying actual behavior.[3]

The first two passages indicate that in devising a grammar for a particular language L, the grammarian is offering a description of one aspect of the speaker's psychological makeup--what Chomsky calls his 'linguistic competence'. The point of the third passage, for our purposes, is that this competence takes the form of an internalized system of rules. Grammar construction, as it is construed here, falls squarely within the boundaries of cognitive psychology. Linguistic functioning is one of the central features of cognition, and the grammarian attempts to describe one very specific aspect of such functioning--the linguisitc rules that comprise the speaker's knowledge of language. This is the dual thrust of Chomsky's mentalism: that linguistic theorizing is about the mind, and that it is about the underlying mental features (as opposed to, for example, overt behavioral features of mind).

All this we may take as 'the psychologization strategy'. What is not yet clear, however, is the precise nature of the relationship between the grammar and the

speaker of the language. We can agree that linguistics is part of psychology, and grant that a generative grammar of the sort Chomsky has proposed is adequate as a description of certain aspects of our linguistic knowledge. Still, it might be argued that the grammarian's concern with competence is like the physicist's concern with bodies in motion. The relationship between the speaker and the grammar could be taken as analogous to the relationship between a boulder and the law of gravity. The law describes and helps predict the motion of the boulder under certain conditions, but the boulder does not know the law of gravity; it simply conforms to it. Similarly, one might hold that the grammar may play a role in description and prediction, but that there are no grounds for saying that the speaker knows the rules of the grammar. It is here, in direct opposition to this view, that Chomsky introduces an epistemological twist. The claim that every speaker has internalized or mastered this system of rules is to be understood as ascribing tacit knowledge of the grammar to the speaker. Specifying exactly what this claim amounts to will be one of our main concerns. In Aspects we are given the following characterization:

> Obviously, every speaker of a language has mastered and internalized a generative grammar that expresses his knowledge of his language. This is not to say that he is aware of the rules of the grammar or even that he can become aware of them, or that his statements about his intuitive knowledge are necessarily accurate.

> Any interesting generative grammar will be dealing with processes that are far beyond the level of actual or even potential consciousness; furthermore, it is quite apparent that a speaker's reports and viewpoints about his behavior and his competence may be in error. Thus a generative grammar attempts to specify what a speaker actually knows, not what he may report about his knowledge. (italics mine)
>
> 4, 5

Particular and universal grammar

So far we have been talking about the relation between a speaker and the grammar of his particular language. It must be noted, however, that linguistic competence operates at a second level--the level of universal grammar--and the epistemological interpretation of competence applies here as well. Apart from the tacit knowledge of a grammar, Chomsky claims that all speakers have tacit knowledge of linguistic universals as well:

> A theory of linguistic structure that aims for explanatory value incorporates an account of linguistic universals, and it attributes tacit knowledge of these universals to the child.[6]

These two claims must be carefully distinguished. A grammar of English, for instance, is a component of a psychological theory that applies to individuals in virtue of their being English speakers. Once we accept this psychologization and see linguistics as part of a larger theory of cognitive functioning, it becomes pertinent to ask not only about the structure of linguistic competence, but also about its development. What is it that

enables an individual to become an English speaker? For that matter, how is it that one comes to master any language?

There is a hidden ambiguity in these questions. We can treat a person's linguistic competence--his tacit knowledge of the grammar--as an evolving epistemic structure (similar in some ways to a person's knowledge of chess). From this perspective we might want to know which rules were learned first, which parts of the grammar are most difficult to acquire, and so on. At the same time, we can concentrate on the features of the speaker that allow for the development of such a structure in the first place. Here our interest is also developmental, but our focus is on the initial state of the language-learner. We assume that the child masters the language on the basis of certain environmental input, and our strategy is to develop a theory that explains how the environmental input--what has been called the 'primary linguistic data'--leads to language mastery. It is this second question that is most relevant to the issue of linguistic universals.

This is the problem of language acquisition that is often presented in terms of Chomsky's engineering metaphor of the black box. A device takes as its input the primary linguistic data that confront the language-lear-

ner, and yields as its output knowledge of language. The problem is to discover the initial internal state of the device on the basis of the given information. For Chomsky, the child's acquisition of the language must be treated in much the same way we treat the adult's mature knowledge. In explaining mature knowledge of language, the idea is to isolate the competence component from the rest of the performance mechanisms that mediate linguistic functioning. Language acquisition will also involve an interlocking system of psychological components--e.g., the perceptual system, memory, etc.. The child's task is to figure out the grammar of the language that confronts him on the basis of the primary linguistic data. Chomsky argues that if the child had no clue, no head start at all, then his task would be impossible. But children do acquire a language--they do internalize a grammar. It follows therefore that they do have some clues. These clues are in the form of tacit knowledge of linguistic universals. The linguist's task is to develop a 'linguistic theory' which is a purported description of the universal principles that the child brings to his language learning task. There is at present no consensus on what types of things--let alone what specific principles--are part of universal grammar. Here we will take it as neutrally as possible: whatever language-specific information the child has at his disposal before the task

of language learning begins. On this construal, universal grammar might include information about the type of grammar that could apply to a natural language, a core grammar of specific rules that hold for every natural language, some set of concepts and categories that are involved in the construction of a grammar, simplicity measures, and so on.

From a Chomskyan point of view, it is the possibility of linguistic theory that makes linguistics worth pursuing. Taken as a psychological theory, a grammar of a particular language has a very narrow and non-natural domain of application--i.e., the speakers of the language in question. We would end up with a different grammar/psychological theory for English speakers, French speakers, and so on. Working on such theories as ends in themselves would be ludicrous; like developing one theory of shillings and pence and another of francs and centimes. The point of investigating the role of money in economy is to eventually emerge with a comprehensive monetary theory that abstracts away from individual differences. So too in the study of language. There are, no doubt, some interesting questions about language that can only be framed at the level of a particular language. Nevertheless, if we are to learn anything general about the mind, we must look at the subject qua speaker rather than as a member of some particular linguistic subgroup.

Although the working linguist's methodological focus is on the particular grammars that differentiate speakers, the ultimate goal is the discovery of the universal grammar that all speakers share.

Armed neutrality: three caveats

The view that speakers have tacit knowledge of a grammar is controversial in itself, but the claim that there is a second level of tacit knowledge of linguistic universals introduces a number of additional problems. As Chomsky sees it, tacit knowledge of a grammar is the product of experience (in the form of the primary linguistic data)[7] and the child's tacit knowledge of the linguistic universals. The tacit knowledge of the universals is not acquired through experience. Rather, it is something innate in the child, something he is born with.[8] It is this innateness hypothesis that links Chomsky's conception of linguistics to the resurgence of traditional philosophical rationalism. In short, if there is such a thing as inborn tacit knowledge of linguistic universals, then perhaps the rationalist picture of mind is correct.

I introduce these issues here to clarify a prevalent confusion. Although the second level of tacit knowledge does involve certain rationalist-sounding claims, the general problems of tacit knowledge and innate knowledge

are completely distinct. Innateness is a genetic question--it concerns the sources of our knowledge; the tacit knowledge hypothesis is neutral on this matter. In terms of philosophical commitments, one might hold that people do have tacit knowledge but that that none of it is innate. Conversely, someone could argue that some of our knowledge is innate but that none of it is buried "far beyond the level of actual or even potential consciousness" in the way Chomsky imagines it to be. As we shall see, the commitment to a rationalist theory of mind can serve as a motivation for trying to establish the claim that speakers have tacit knowledge of linguistic universals; a motivation, however, is not an argument.[9]

My primary concern in this thesis is with tacit knowledge, not innateness. For now, I am more interested in the relation between knowledge and consciousness than in the sources of our knowledge. Therefore, I will be defending the view that speakers have tacit knowledge of a grammar and, as far as possible, I will sidestep questions about linguistic universals. Since my focus is Chomsky's theory of grammar I will not be able to steer totally clear of the rationalist position that he defends. For the most part, however, I will have very little to add to this already entangled issue.

Since I am issuing caveats, let me add two more. Throughout this paper I will be speaking of tacit knowledge of a grammar. Because of Chomsky's tremendous influence in linguistics, the term 'grammar' is often automatically associated with his detailed conceptions of what a grammar contains, how it is organized, etc.. Although the original claim that speakers have tacit knowledge is due to Chomsky, one can adopt this view without making any specific commitment as to the nature of the grammar. I have no argument to offer here against Chomsky's linguistic views, but I do want to avoid any such commitment. There is no reason to yoke the fortunes of the epistemological notion of tacit knowledge to the latest views on the constitution of a grammar. My view here is that there is a set of linguistic rules that speakers tacitly know. The specific nature of these rules is an independent question.

My last caveat takes the second a step further. I have just claimed that my position on tacit knowledge is neutral as regards the specific structures of grammars. Now I want to consider the following objection. Suppose someone argues that although I am not committed to any specific view of grammar construction, I am committed to Chomsky's conception of linguistics as a branch of cognitive psychology. The view that speakers have tacit knowledge of a grammar is, after all, a development of

just this line of thought. What if this conception of linguistics is undermined and rejected? What of a philosopher who claims that a grammar is a formal demarcation of "the sequences capable of occurring in the normal stream of speech"?[10] Or the claim that a grammar is not a psychological theory but a theory of a language as an abstract entity? I am not interested here in the viability or fruitfulness of these conceptions of linguistics and grammar; we will touch on some of these issues briefly at a later point.[11] My only point is that although these alternative conceptions do not take a grammar to be a theory of an underlying psychological reality, they are still not incompatible with what I want to claim about speakers' tacit knowledge. The case for tacit knowledge, then, is <u>not</u> conceptually bound to Chomsky's psychologization of linguistics. The tacitly known rules might not conform to these non-Chomskyan conceptions of what the/a grammar of a language is. The set of rules might fall short in terms of simplicity, completeness, or whatever. Nevertheless, whatever one thinks a grammar is about, there is no denying the fact that most humans have mastered a natural language. My thesis is that such mastery is to be explained in terms of an underlying tacit knowledge of linguistic rules.

<u>Conditions of Adequacy</u>

But now we have to face a question of strategy. How does one go about defending the idea that there is such a thing as tacit knowledge? One is tempted to say that Chomsky's comments are suggestive, but that they do not give us what we need--i.e. a theory of tacit knowledge. There is a persistent feeling that if we satisfied this magical requirement--if we provided a theory--then the concept of tacit knowledge would be instantly embraced. But can such a theory be provided? Certain key elements seem to be lacking. We can set out to construct a theory of perception, for instance,because we have an intuitive sense of what perceptual phenomena are. In a less rigorous sense of the term, philosophers talk about a theory of knowledge or belief as whatever it is that will provide solutions to the age-old puzzles and questions about these notions. Our case, however, is different. It should be obvious that there is no pre-theoretically well-defined set of 'tacit knowledge phenomena' to be explained or predicted. There are no classical problems of tacit knowledge that have troubled philosophers through the ages.

In the light of all this, I want to suggest that a defense of the notion of tacit knowledge involves legitimation rather than theory construction. We must establish the philosophical credentials of the concept by showing first, that it is conceptually coherent, and

second, that it plays a significant role in a broader explanatory theory. There is nothing new here; this dual requirement obtains whenever we introduce a new theoretical concept.

In the case at hand, the first requirement translates straightforwardly into the question 'Is the concept of knowledge that is not even potentially conscious internally consistent, or is it a theoretical impossibility?' This is the topic of the next chapter. The second requirement is a little more vague. In order to sharpen it I want, in what remains of this chapter, to consider a set of conditions of adequacy on the introduction of unconscious mental states in general, and then adapt these conditions to the case of tacit knowledge. These conditions are derived from Brentano.

In Psychology from an Empirical Standpoint Brentano asks whether ". . . all mental phenomena are conscious, or might there also be unconscious mental acts?" Although he ultimately rejects the idea that there are any unconscious mental phenomena, his argument is empirical. He finds nothing conceptually amiss with the opposing view; it is simply not supported by the evidence. In the course of his discussion he advances a set of empirical conditions which, he contends, must be satisfied if we are to accept the existence of unconscious mental

phenomena into psychology. For us, the relevant members of this set are as follows:

> (1) In order to be able to draw any conclusion concerning an unconscious mental phenomenon as a cause, from a fact which is supposed to be its effect, it is necessary, first of all, that the fact itself should be sufficiently established.
>
> (2) . . . it is necessary. . . to show through experience that conscious mental phenomena have always involved similar consequences.
>
> (3) . . . it is necessary to assume that they (conscious mental phenomena) do not involve at the same time other consequences which are absent at the case at hand, even though there is no reason to suspect that they are connected with the concomitant consciousness which is missing in these cases.
>
> (4) . . .it is necessary that the unconscious mental phenomena, to which the hypothesis appeals, do not contradict, in their succession or in their other characteristics, the recognized laws of conscious mental phenomena, so that any possible peculiarities can be understood on the basis of the lack of concomitant consciousness,. . ..
>
> (5) . . .the origin of the mental phenomena assumed to exist despite the absence of consciousness should not be considered to be something utterly and entirely inconceivable in itself.
>
> (6). . . the phenomena under discussion cannot be understood , at least not without the greatest improbability, on the basis of other hypotheses.[12]

It will be convenient to break down this set of conditions into two groups. Although they are all framed in terms of unconscious mental phenomena, (1), (5), and

(6) are really general constraints on all theoretical posits. (2), (3), and (4) on the other hand, are more specific in their application. Let us start with the first group.

(1) is simply a warning: if our grounds for introducing an unconscious mental state s is its role in the causal explanation of some state or event e, then we had better be sure that there is an e which we've described accurately, and which would be explained if s existed. This warning is not idle in our case as there is some controversy regarding what it is that ascriptions of tacit knowledge are supposed to explain. Shall we take the explanans to be the ability to speak and understand the language? Or is it the fact that speakers have intuitions about the grammatical properties of expressions in their language? Perhaps we should concentrate on the speaker's ability to notice and correct linguistic errors? These are distinct, and as Hume would say, separable, linguistic phenomena. We will have to specify which, if any, require ascription of tacit knowledge for their explanation.

(5) adds the constraint that we be prepared to say something about the source of the proposed unconscious mental phenomena. This is relevant in two ways in the tacit knowledge case. First, it has been argued that if

we consider the proposed sources of the speaker's tacit knowledge of a grammar, the theory collapses into an infinite regress. Looking ahead a bit, we find one form of this argument which proceeds roughly as follows:

> To tacitly know a grammar one must master some sort of language in which to represent the grammar. But since our theory implies that one must tacitly know a grammar before one can master any language, we are saddled with an infinite regress of grammars and languages.

The solutions to this puzzle and others like it will be discussed later.[13] Here I simply want to point out that a theory without an answer to this regress-challenge would fail to meet condition (5). A second sense in which (5) could be relevant to our concerns involves the innateness hypothesis discussed earlier. If one chooses to explain the speaker's tacit knowledge of grammar in terms of his tacit knowledge of linguistic universals, then one must be prepared to show that this hypothesis is at least better than an "utterly and entirely inconceivable" explanation of how the speaker comes to have tacit knowledge of a grammar.

We cannot legitimate the notion of tacit knowledge by simply showing that it <u>can</u> be given a role in an explanatory theory; we have to show that it <u>must</u> be given such a role. To satisfy condition (6) we must show that no other available hypothesis can explain the phenomena in question and is simpler than ours. (6) is a principle

of last resort; we don't advert to explanation by unobservables unless all attempts at explanation in terms of observables fail. A good portion of our defense, therefore, will involve looking at alternative explanatory models that attempt to account for all the relevant phenomena without employing the notion of tacit knowledge. So much for the first group.

The second, third, and fourth conditions are all instances of what we can call the Principle of Analogical Dependency. The idea behind this principle is not very complicated. In introducing the idea of an unconscious mental state--take unconscious desires as an example--we want to sever some of the connections that hold between the concept of desire and the rest of our conceptual scheme. For instance, we are willing to suspend some of the ordinary behavioral criteria for ascribing desires if we have unconscious desires in mind. However, if such hypothesized desires did not interact with the subject's beliefs to affect his behavior, then we might begin to wonder. Perhaps unconscious desires are really not desires at all; no more than alleged criminals are criminals. If we sever too many connections, then it seems that we are no longer talking about desire. The term 'unconscious desire' loses its explanatory power and becomes a misleading metaphor.

The lesson here is that talk of an unconscious mental state is in some sense parasitic upon talk of its conscious counterpart. The three conditions in the second group can be seen as attempts to spell out this dependency. As adapted to our case, (2) demands that explanations involving tacit knowledge run parallel to explanations involving explicit knowledge. In other words: if we want to explain some state or event e in terms of an individual's tacit knowledge, we first have to show that knowledge-based explanations are appropriate for e-type phenomena. This specific demand is at the heart of a Ryle-inspired objection that we will consider later--i.e., knowledge of language is best analyzed as an ability, and abilities are not the kinds of things that are to be explained in terms of underlying knowledge.[14] (It follows a fortiori that abilities are not to be explained in terms of underlying tacit knowledge.)

The third and fourth conditions of adequacy are very close in spirit. (3) is a fairly simple requirement, but it is not easy to state precisely. The general idea is that we must be able to explain any disanalogy between a conscious mental state and its unconscious counterpart as a function of the conscious awareness that is present in the former case and absent in the latter. A possible example might be as follows. If a linguist has conscious knowledge of a grammar of a language, he might be able to

provide information about the transformational history of a sentence in that language. A native who, by hypothesis, has tacit knowledge of a grammar does not have this ability. (3) requires that the linguist's special ability be explainable in terms of the fact that his knowledge is explicit. Overall, we will have to show that there is no reason to think that there are disanalogies that are not amenable to such explanation.

(4), when relativized to our case, requires that tacit knowledge be genuine, full-blooded knowledge; that it conform to the best theory of knowledge we have available. If we accept a justified-true-belief analysis of knowledge, then we will have to show that tacit knowledge satisfies these conditions; similar considerations apply if we accept a causal theory.

NOTES CHAPTER I

1. I don't mean to imply that there is nothing to be learned from examining the logic of the skeptic's position--only that it is no longer a real threat as it was to Descartes. For more on this first point see Gilbert Harman, Thought (Princeton, NJ: Princeton University Press, 1973), ch. 1.

2. For the sake of brevity I will refer to speaker-hearers as simply 'speakers'.

3. Noam Chomsky, Aspects of A Theory of Syntax (Cambridge, Mass.: The MIT Press, 1965), p. 4.

4. Ibid., p. 8.

5. To avoid confusion I want to clarify what I think Chomsky means by the phrase 'beyond the level of actual or even potential consciousness'. Imagine that Jones' memories of the war are so traumatic that his defense mechanisms make it practically impossible for him to recall them; he would have a nervous breakdown first. The memories may be said to be beyond even potential consciousness in the sense that, given our confidence in the background theory, we can confidently predict that the subject will never recall these experiences. Although there are critical differences between the two contexts, I think this is close to the sense that Chomsky intends here.

There are other senses, however, which are much stronger. One might say that the experience of finding your way around your living room the way a bat does is beyond the level of even potential human consciousness. Here the state or process in question is of the wrong type. Grammatical rules are obviously not of the wrong type. If the linguist were to explain the rule to the subject, the subject would presumably grasp what was being communicated and in this sense would be conscious of it. What he would not be aware of, if Chomsky is right, is that these rules are the rules he has been applying.

6. Chomsky, Aspects, p. 27.

7. How this notion of a 'product' should be explicated will be discussed in later chapters.

8. More accurately: it is part of the child's genetic heritage. See Stephen Stitch, ed. Innate Ideas (Berkeley: University of California Press, 1975), "Introduc-

tion".

9. The rationalist theory of mind is explicitly cited as a motivation in C. Graves, J.J. Katz, et. al., "Tacit Knowledge" in JP, 11, June 7, 1973.

10. Here I have Quine in mind; see note 67 of Chapter II for the relevant biographical data.

11. See pp. 58ff.

12. Franz Brentano, Psychology from an Empirical Standpoint, English ed. L. McAlister, (NY: Humanities Press, 1973), pp. 106-109.

13. See Chapter III.

14. See Chapter III, Section B.

CHAPTER II

TACIT KNOWLEDGE: SOME CONCEPTUAL OBJECTIONS

In this chapter I want to consider a number of a priori or conceptual criticisms of Chomsky's ascription of tacit knowledge of a grammar to speakers. I take these criticisms to be a priori in that they attempt to show that the notion of tacit knowledge is conceptually incoherent; that all knowledge must be explicit. If this claim could be established, then the epistemological twist of Aspects which we alluded to earlier would have to be rejected. It is therefore imperative that we deal with this challenge at the outset.

The criticisms I will take up fall into two groups. The first set objects to tacit knowledge from almost an ordinary language point of view--i.e. the term 'knowledge' is only applicable within a set range, and what Chomsky refers to as 'tacit knowledge' is simply outside that range. The second group contends that any attempt to explain behavior in terms of tacit knowledge ascriptions is doomed to failure for general methodological reasons.

Belief Ascription and Its Linguistic Expression

We may begin with the familiar fact that, on most views, knowledge entails belief. It follows then that if tacit knowledge is a genuine epistemological category, we will have to consider ascriptions of tacit knowledge as (at least) involving ascriptions of belief. This is no idle requirement; as we noted in our discussion of Brentano's conditions, it means that such ascriptions must satisfy the constraints we impose on standard ascriptions of belief. It will be pointed out that belief is closely connected to linguistic expression of belief; a constraint on standard ascriptions is that the subject be able to tell us what he believes. This paradigmatic requirement will carry over *ipso facto* to ascriptions of knowledge as well.

Belief and avowal

But now it seems that we have a problem: it is just this verbalizability that is absent in Chomsky's talk of tacit knowledge of a grammar. In fact, Chomsky is quite explicit in rejecting this requirement--we recall his claim that the knowledge in question is "below the level of actual or even potential consciousness." If the subject cannot even become conscious of these beliefs, we certainly cannot expect him to state them. The argument, simply put, is that by rejecting this paradigmatic constraint on belief-ascription, Chomsky makes it impossible

for us to take the notion of tacit knowledge literally. Of course there might be some interesting relation between the rules of a language and a speaker, but it cannot be anything like knowledge.

To assess this argument, we have to clarify at least two points. First, exactly what sort of linguistic behavior does our critic take to be necessary for belief-ascription? Must the subject be able to state what we impute to him as something he believes or will some less explicit sort of verbal expression suffice? Second, what is the nature of the (purported) relationship between ascription of belief and its linguistic expression? Is the latter a necessary condition for the former? Criteriologically connected? Analytically tied? Let us take these questions in order.

The common-sense intuition that supports the view that belief and linguistic expression of belief are closely connected is this: if A is claimed to know that p, but A cannot give a statement of p (or at least come close to stating it), then ordinarily we do not credit the claim that he knows p. This is the standard test we apply in verifying third-person ascriptions and it would seem to hold for first-person avowal as well. This common-sense intuition, if taken literally, has very little to recommend it. People certainly believe things even

though they are in no position to state what it is they believe. Sleep, paralysis, aphasia, and tip-of-the-tongue phenomena are but a few examples of the sorts of things that would provide conclusive counterexamples to this view. It is an undeniable fact that belief does not always manifest itself in avowal.

Recognition and inner recognition

This radical position is clearly unacceptable; it is therefore no surprise that nobody holds it.[1] A less radical position would be that if the subject is able to recognize some formulation of what he believes, then we can ascribe the belief to him nonetheless. That this recognition-requirement might be a necessary condition for belief-ascription is suggested by the following example. Consider a frustrated game-show contestant who claims to know who discovered the Pacific Ocean even though he cannot think of his name at the moment. After his time is up, imagine that the moderator tells him that the correct answer is "Balboa". If the contestant looks puzzled and fails to acknowledge "Balboa" as the correct answer, then we are apt to say that he was simply confused in thinking that he knew the answer. But if he claims to have known it all along, and evinces the proper frustration--pulling his hair, groaning loudly, etc.--we might agree that he knew it all along. We could then go on to attribute his failure to the time-pressure, to a

temporary lapse of memory, or whatever. Perhaps, then, what we should demand as a constraint on belief/knowledge ascription is not full-fledged avowal but rather recognition: if A knows/believes p then he will recognize some formulation of p as an expression of what he knows/believes.

Some initial difficulties. This condition is certainly more reasonable than the avowal requirement taken up earlier, but it still has its problems. In the first place, we have a right to be skeptical about the possibility of formulating an explicit criterion for recognition. An example should show why. Let's say that I have just been taught something about time-signatures in contemporary music, and I am told that John Coltrane's version of "My Favorite Things" is played in 3/4 time while most other versions are not. Since I am familiar with the performances in question I might say that I knew all along that the Coltrane version had a different time than the original, but I just didn't have the words to express this knowledge. Are we to say that I really knew or believed all along that Coltrane's version was in 3/4 time? Does this sort of thing count as recognition? It is hard to say.[2]

Apart from this difficulty, which seems to be tied up with a systematic vagueness of the term 'recognize', it has been suggested that there is another deeper problem. One typical case of belief were the avowal condition is too strong is psychoanalytic ascriptions of unconscious beliefs. A person going to a psychiatrist might be informed that he has a tremendous inferiority complex because of his weight; that he believes people think less of him because he is fat. The patient might very well deny that he has any such belief and adduce what he takes to be counterevidence against the therapist's claim. After a while, though, we can imagine the psychiatrist breaking through: how else can the patient account for the unhealthy and uncomfortable girdle he wears? for his proclivity to choose friends of similar or greater bulk? for his tendency to lie when asked his weight? The doctor might point out that even though the patient consciously thinks of himself as a loved and lovable fat man, in ninety-nine percent of the cases in which patients displayed similar 'symptoms' it turned out that they had inferiority complexes associated with weight. The patient, confronted with the mass of evidence the therapist presents, eventually agrees: "Well, I guess you must be right. I must have the inferiority complex you told me about and I've probably believed all along that people undervalue me because of my girth.

What else could explain the symptoms you point out?"[3]

Nagel's condition and psychoanalytic ascriptions. In one colloquial sense of 'recognize' we can say that the patient here finally recognizes that he is suffering from an inferiority complex. There is another sense of the term, however, in which we must say that he does not truly recognize the complex. We might say that he has discovered, with the aid of his therapist, that he suffers from the complex, but there is no real recognition. An explanation of what is missing is offered by Nagel.

> The psychoanalytic ascription of unconscious knowledge, or unconscious motives for that matter, does not depend simply on the possibility of organizing the subject's responses and actions in conformity with the alleged unconscious material. In addition, although he does not formulate his unconscious knowledge or attitude of his own accord, and may deny it upon being asked, it is usually possible to bring him by analytical techniques to see that the statement in question expresses something that he knows or feels. That is, he is able eventually to acknowledge the statement as an expression of his own belief, if it is presented clearly enough and in the right circumstances. Thus what was unconscious can be brought, at least partly, to consciousness. It is essential that his acknowledgement not be based merely on the observation of his own responses and behavior, and that he come to recognize the rightness of the attribution from the inside.[4, 5]

The imaginary patient we have been discussing does not come to see from the inside--there is no "Aha! So that's what's been bothering me all along; I knew it."

All we have is the patient reacting like any other objective observer. The evidence is presented, the statistical regularities are noted, and the rational man simply agrees that the case has been made: whoever is suffering these symptoms is suffering from such and such a complex. That the 'whoever' happens to include the patient himself is incidental. In this sort of case Nagel wants to say that the ascription of the unconscious belief (that others devalue the patient because of his weight) is premature. In order to ascribe unconscious belief to the patient, the therapist must believe not only that the patient as rational man can be convinced that he has such a belief, but also that the patient-as-complex-sufferer can be brought to the aha-experience, i.e., that the patient has the potential to acknowledge the belief as his own, 'from the inside'.

Nagel generalizes this feature of the psychoanalytic case into a necessary condition on all ascriptions of epistemic states:

> It seems to me that where recognition of this sort is possible, there is good reason to speak of knowledge and belief, even in cases where the relevant principles or statements have not yet been consciously acknowledged, or even in cases where they will never be explicitly formulated.[6]

He also seems to hold the converse: that where such recognition is not thought possible, we should not ascribe

such epistemic states. In assessing Chomsky's claim that speakers have tacit knowledge of a grammar, Nagel goes on to suggest that

> accurate formulations of grammatical rules often evoke the same sense of recognition from speakers who have been conforming to them for years that is evoked by the explicit formulation of repressed material that has been influencing one's behavior for years.[7]

Where such recognition is possible we may ascribe tacit knowledge of the rules to the speaker. Where it is not we can only say that the suggested rules 'describe his competence' but not that they are followed.

It is hard to know what to make of Nagel's position. On the one hand, the arguments put forward seem, at first, at odds with the claim that speakers know a grammar tacitly. Nevertheless, the conclusion reached--that speakers <u>do</u> have knowledge of grammar--is more sympathetic than antagonistic. My own view is that Nagel is sympathetic for the wrong reasons; his premisses are better used (and have been so used) in making the opposing argument. Ultimately, I think that the argument that emerges is, at best, inconclusive. But first I want to say something more about Nagel's position.

Nagel's misplaced sympathies. We should mention, first off, that Nagel is assuming that the ascription of tacit belief in the psychoanalytic case is independently

justified, or at least within the bounds of the paradigmatic use of the terms 'belief' and 'knowledge'. If this were not so, there would be very little point in assimilating one controversial context--tacit knowledge of the rules of a grammar--to another controversial one--unconscious belief and knowledge in psychoanalysis. Here I agree with Nagel in accepting the psychoanalytic ascription context as legitimate. Still, as we shall see later, this acceptance is by no means universal.

I want to explain now why I take Nagel to be sympathetic for the wrong reasons. The crux of it is his peculiar view that the rules of a grammar have the same sort of recognition potential as repressed material in the psychoanalytic setting. From a linguistic point of view this is naive. As the experience of anyone who has studied syntax will show, recognizing a rule 'from the inside' is an extremely rare (if not nonexistent) occurrence. Of course it is always possible that the best theory of English syntax available at this point is not at all like the optimal theory of English that we strive for. It could be that the optimal theory will be made up exclusively of rules that will be recognizable when we present them to speakers (assuming, of course, that the speakers are properly informed about the theoretical vocabulary and so on). Even if we grant this

rather far-fetched possibility, however, we are left with a rather odd methodological perspective on linguistic research: any proposed rule of a grammar is in principle testable by exposing those who have the appropriate theoretical background and vocabulary to the rule. If such an individual--i.e., a linguist--recognizes the rule (from the inside) as an expression of what he has been following all along, then we can count it part of his tacit knowledge; otherwise, we cannot. Each linguist is in this regard a theoretical oracle. This position has too many problematic consequences to consider them all here. To take but one example, consider the possibility that one linguist recognizes a candidate rule and another does not. The hope for developing a theory along these lines seems very dim.

The basis of Nagel's naivete about recognizability potential is his construal of grammatical rules. He takes as an example "that the plural form of a noun is usually formed by adding 's', and that among the exceptions to this is the word 'man', whose plural is 'men'."[8] I think it is an open question whether the normal speaker would recognize even such a rule 'from the inside' or perhaps as only a blatantly obvious regularity that his behavior conforms to. Whatever the case, this morpho-phonemic generalization is in no wise representa-

tive of the sorts of rules that linguists now think comprise a grammar. As I indicated above, the level of complexity of grammatical rules is far beyond that of the plural rule that Nagel alludes to. An example of the former sort of rule might be this:

Tag Formation (Optional)

SD: Q - (not) - NP - Tense $\begin{pmatrix}\text{Modal}\\ \text{have}\\ \text{be}\end{pmatrix}$ - X

1 2 3 4 5

2 3 4 5 1 3 4
+PRO

It is hard to imagine that anyone would claim that such a rule, which is involved in the derivation of (2) from (1), could

(1) John will go

(2) John will go, won't he?

evoke the recognition from the inside that Nagel demands.[9] If linguists are on the right track, then acceptance of Nagel's aha!-condition should lead to the rejection of the view that speakers have tacit knowledge of the grammar of their language.

Cooper's argument. This, in fact, is the position adopted by Cooper as part of his attack on Chomsky's tacit knowledge claims. He agrees with our argument that the actual rules of a grammar are best thought of as

beyond the aha!-sort of recognition experience, but also agrees with Nagel that such recognition <u>is</u> a necessary condition for belief-ascription:

> Since statements of the rules, structures, etc. of a generative grammar typically elicit no such reaction, the condition for attributing knowledge of these rules or structures is missing. [10]

So although we have found Nagel's conclusion wanting, his premisses can be used to argue that ascription of tacit knowledge of a grammar is unwarranted. What then are we to say about the premisses? Is there any argument to be made in support of the view that belief-ascription and the expression/recognition of belief are linked up the way Nagel thinks they are?

Nagel makes no explicit argument but one senses that for him the analogy with the ascription of unconscious beliefs in the psychoanalytic setting must carry the weight. After all, when we think about a belief as tacit we are, more or less, claiming that although it is not conscious, it is still 'in the mind'. Whatever the differences in content or in practical consequences, Freud's unconscious beliefs have this characteristic as well. It is therefore quite natural to suggest that the conditions we deem necessary for ascription in the Freudian case apply to the linguistic case as well.

Freud and inner recognition. But appealing to Freud will not help. I will argue that we have no sound reason to take Nagel's recognition from the inside to be a condition of unconscious belief ascription even in the psychoanalytic case. I think that Nagel has misconstrued the Freudian view as to the relation between belief and the aha!-experience.[11] For Freud, such recognition on the part of the patient--actual or potential--is never a condition on the therapist's ascription of unconscious beliefs, motives, etc.. (Freud holds that all these may be unconscious in the required sense.) Perhaps it is a requirement for successful therapy and rehabilitation, but that is an entirely different matter. There seems to me to be strong evidence that for Freud, the ascription of unconscious mental 'conceptions' was strictly a matter of empirical inference to the best explanation. In two essays devoted to metatheoretical issues involving the unconscious, he never once indicates that inner recognition is necessary for ascription. He does, however, give two sorts of arguments, one negative and one positive, for accepting the unconscious. I want to discuss these arguments because aside from clarifying the relation between belief and ascription in the Freudian psychoanalytic setting, they seem to me to be applicable to the ascription of tacit knowledge of a grammar as well.

Freud and the unconscious. Freud's negative argument is just that there is no argument against the unconscious; to deny that there are unconscious mental events on conceptual grounds "...simply begs the question by asserting 'consciousness' to be an identical term with 'mental'."[12] And again, that such an identification

> is either a petitio principii and begs the question whether all that is mental is also necessarily conscious, or else it is a matter of convention, of nomenclature. In this latter case it is of course no more open to refutation than any other convention. The only question that remains is whether it proves so useful that we must needs adopt it. To this we may reply that the conventional identification of the mental with the conscious is thoroughly unpractical.[13]

The positive argument is no more and no less than Freud's theory of mind. In other words, the ascription of unconscious beliefs is more practical in terms of explanation of mental phenomena than the denial of such ascriptions. This argument is incorporated in his claim that

> All these conscious acts remain disconnected and unintelligible if we are determined to hold fast to the claim that every single mental act performed within us must be consciously experienced; on the other hand they fall into a demonstrable connection if we interpolate the unconscious acts that we infer. A gain in meaning and connection, however, is a perfectly justifiable motive, one which may well carry us beyond the limitations of direct experience. (Italics mine)[14]

As the italicized section should indicate, there is no

'favored evidence status' accorded to recognition, conscious apprehension, etc..[15] Freud takes this claim a step further by explicitly severing the conceptual connection between ascription of unconscious states and processes and the therapeutic possibility of bringing the patient to acknowledge the belief from the inside:

> Even if psychoanalysis showed itself as unsuccessful with all other forms of nervous and mental disease as with delusions, it would still remain justified as an irreplaceable instrument of scientific research.[16]

Theory and therapy. As I said earlier, one way of viewing Nagel's mistake is to see it as a result of his confusing what might be required for <u>therapy</u> with the conditions for construction of a <u>theory</u>. It might be that successful treatment of psychological disorders is possible only if the patient comes to recognize his unconscious beliefs, desires, etc., 'from the inside'--this is feasible only to the extent that the patient's 'resistance' can be broken down. Nevertheless, even if the therapist becomes convinced that the resistance cannot be broken down--that the patient, for whatever reason, cannot be brought to recognize from the inside--he does not thereby abandon his diagnosis. Quite the opposite: if the diagnosis takes account of the patient's overwhelming resistance, then the therapist's theory (of this specific patient) will ascribe an unconscious belief that the

subject cannot be brought to acknowledge from the inside.

One might be tempted to say that the foregoing discussion of the Freudian position is only of historical interest. Even if this charge were true, we should be consoled by the fact that, as Flew says, "Freud's stature was sufficient to justify the study of his ideas in their own right."[17] I don't think it is true and we shall see why in what follows.

Belief and expressibility: three views

The lesson of this short digression is this: there is no basis in Freudian psychoanalytic theory for the claim that recognition 'from the inside' is a necessary condition for belief ascription. But knocking out the Freudian foundation for Nagel's claim still leaves open the possibility that there is another argument that can serve to establish the connection. I know of no way to show that there can be no such argument, but we can go one step further by considering the general forms such an argument might take. I therefore want to turn away from the Freudian view for a moment and look at the problem more broadly. It seems to me that if one were to attempt to establish a connection between belief-ascription and recognition, one would have to hold that the connection is either (a) analytic, (b) criterial/paradigmatic, or (c) evidential. I want to show that there is no case to

be made along any one of these three lines that raises problems for ascription of tacit knowledge of a grammar.

The Analytic Connection. One might argue (although I know of no philosopher who does so explicitly) that it is part of the meaning of 'A believes that p' that A can express p linguistically or at least be brought to recognize p from the inside. Then, since speakers can not in general recognize p where p is a standard formulation of a grammatical rule, the claim that speakers have tacit knowledge of a grammar is unwarranted--it is without meaning. At best, we would have to take 'tacitly knows' or 'tacitly believes' as expressing a new relation having very little to do with ordinary knowledge and belief.

Although not a 'friend of analyticity', Goodman at one point argues as if something at least as strong as Nagel's condition of recognizability is conceptually tied to the attribution of ideas. Whether he would use the same argument against the notion of tacit knowledge is not clear, although the following passage does occur as part of a critique of Chomsky's a attempt to revive the rationalist theory of innate ideas.

> A: And being ideas, they are in consciousness?
>
> J: No, not necessarily; not even usually.
>
> A:Then they are in the subconscious mind, operating upon cognitive processes, and capable of being brought into full consciousness?

> J:Not even that. I may have no direct access to them at all. My only way of discovering them in my own mind may be by the same methods that someone else might use to infer that I have them, or I to infer that he does.
>
> A: Then I am puzzled. You seem to be saying that these innate ideas are neither innate nor ideas. (Italics mine) [18]

This sort of view is reminiscent, perhaps, of an early logical behaviorist conception of belief.[19] There is a touch of irony here as well. One of the driving motivations of the logical behaviorists was to free ascription and talk of belief, etc. from the internal baggage it had been forced to carry for so long. Belief, on this liberated view, is totally a matter of publicly observable behavior.[20] It is strange, then, that these philosophers would revert to the view that there is one favored sort of behavior--i.e., recognition, verbal expression, and so on--that is in some sense equivalent to belief. Belief is now not to be identified with any single internal state, but with a single behavioral disposition.

The general arguments against this view of belief (behaviorist/dispositional) are too well-catalogued and need no repetition here. This narrower view under consideration--that the connection between belief-ascription and behavior is analytic and that certain types of dispositions to verbal behavior deserve special status--

faces additional problems. First, I can find no real argument in its favor. Second, the thesis is at odds with our ordinary application of the terms 'belief', 'knowledge', etc., to organisms that are not capable of the special sort of linguistic expression claimed to be necessary. I do not want to get embroiled in this complex question here except to make the following two points. First, the burden of proof is still on those who would argue that ascription of belief to dumb animals and prelinguistic infants is unwarranted or even that animals and infants 'believe' in only a 'logically secondary' sense.[21] Second, even if we were to grant such claims, they do not threaten our position on the speaker's tacit knowledge of a grammar. In ascribing such knowledge, we are certainly talking about organisms that do have language; no one doubts that speakers have a great many beliefs that they are able to express linguistically. So ultimately it will not be enough to argue that language (in the sense of 'being a language user') is central to belief (in the sense of 'being a believer'). What needs to be shown is that this connection is essential--in fact definitional--in the case of every single individual belief; that for every p, if A believes p, then A can at least recognize a linguistic expression of his belief. This, to my knowledge, has not been shown.[22]

The Criterial/Paradigmatic Connection. Another possible move is to claim that the relation between belief ascription and recognition from the inside is not an analytic one but something weaker. It is not part of the meaning of 'A believes that p' that A have any specific linguistic disposition; rather, we might say that one of the minimal 'criteria' for ascribing beliefs involves recognition of the appropriate sort. We might say that it is part of the 'paradigm' use of the notion of belief that believers have at least this sort of disposition. I lump these two (perhaps different) conceptions of the relationship together because neither is clearly understood. If there are such relations--i.e., 'criterial' or 'paradigmatic'--they stand, very conveniently, in the twilight zone between conceptual necessary connections and contingent evidential ones. This position, unlike the first, does have at least one adherent; as I indicated earlier, it is explicitly defended by Cooper.

> To justify extension of terms like 'knowledge' and 'belief' beyond the paradigms, there must be sufficient analogy between the extended cases and the paradigms. There is a fairly smooth progression from being able, here and now, to avow knowledge of some rule R, through being unable to do so only because one has momentarily forgotten the rule, to the case of being ready to acknowledge 'R', once it is formulated for the person, as a description of a rule with which one's behavior has accorded... Merely behaving in a way that a person who paradigmatically knows the rule might behave provides too tenuous a link with the paradigms to warrant extension of the relevant

terms.[23]

We should note at the start that although Cooper talks of paradigm cases and extensions of paradigms, he never puts forward a general theory about how such extensions are to be constrained. Instead, he seems to assume that there is an accepted theory that all philosophers subscribe to. If there is such a theory, I have never some across it.[24] It is important to realize that the issue is not over the paradigmatic ascriptions of knowledge and belief, but over the question of when extension of terms to non-paradigmatic cases is legitimate. To argue that the term 'knowledge' or 'belief' cannot be stretched to cover the new context of tacit knowledge of grammatical rules, one must appeal to a set of general principles that govern the flexibility of these terms. Without such accepted principles, there is a gap in the argument. This gap I take to be serious, and not adequately filled by Cooper's remarks about 'smooth progressions'. I will argue that the limits of the applicability of terms like 'knowledge' and 'belief' can not be determined by simply searching out smooth progressions. To the extent that we take the notion of a paradigm-use seriously, the ascription of tacit knowledge of grammatical rules to speakers is a legitimate extension of the paradigm. Let us look at these points in order.

The concept of a 'smooth progression' is something like Wittgenstein's conception of a family resemblance. If we don't specify what the progression is to be in terms of, we can expect to find that any two contexts can be smoothly linked. The specific contexts which Cooper picks out, (1) immediate avowal, (2) avowal after 'memory search', (3) acknowledgement,[25] are clearly points along a line. What he fails to consider is that the progression is only 'smooth' in regard to levels of the relation of belief and its linguistic expression. But no one has ever claimed that tacit knowledge of grammatical rules (as we've conceived of it following Chomsky) belongs anywhere on this progression; this sort of tacit knowledge is admittedly not a case of verbally expressed belief/knowledge. What we are trying to show, rather, is that it is a case of belief/knowledge simpliciter. To assume that the stages from immediate avowal to acknowledgement 'from the inside' exhaust the contexts for the proper application of the notions of belief and knowledge is to beg the very question at issue. It already embodies the assumption that all knowledge/belief involves a present capacity for linguistic expression or recognition.

We might also notice that Cooper's dividing line between the paradigmatic applications and the extended cases is not where it belongs. If we think of the ex-

tended cases as those contexts where application of the given term involves some 'stretching'--i.e., where we are dealing with an 'extended sense' of the term--then it is hard to see why all the stages he mentions are not part of the paradigmatic use. Perhaps we are dealing with an ambiguity here. In one sense, the 'paradigm case' of a concept is the standard application or stereotypic context; one which we might advert to in explaining the use of the term to someone who does not yet understand it. In another sense, the possibility of extending a paradigm to cover a new case indicates not only that it is not the standard, classic instance, but also that we create an 'extended sense' of the term when we use it so. Consider the following example. A paradigmatic chair (if there is such a thing) is a straightforward piece of furniture that looks something like the picture we might expect to find in a dictionary, and this I think corresponds to the first sense of 'paradigm case'. Now consider some novel configuration of plastic tubing and canvass that, as a surprise to everyone but the designer and manufacturer, is also a chair.[26] Perhaps we wouldn't want to call this a 'paradigm case' of the concept chair; nevertheless, this contraption is no less a chair than the Shaker chair in the catalogue or dictionary. We don't use 'chair' in any extended sense when we apply it to this new model. Now when Cooper talks about acknowledging-from-the-inside

as two steps away along the smooth progression from the paradigm of having a belief or following a rule--i.e., immediate avowal--he could mean one of two things. He might hold that in the case where there is no immediate recognition we have only an extended sense of 'knowing Rule R'--i.e., that even where Nagel's condition _is_ satisfied we don't have real knowledge or belief. The other possibility is that acknowledging from the inside is not _paradigmatic_ only in the sense that it is not _stereotypic_; still, where Nagel's condition is satisfied, we _do_ have bona fide knowledge/belief. On this issue, Cooper is quite explicit:

> Avowal is a criterion for paradigmatic ascriptions of knowledge and belief. Readiness to acknowledge 'from the inside' is a criterion for ascribing knowledge and belief in other cases. And there should be no temptation to think that the paradigm and extended cases are identical except for the alleged evidential symptoms. The difference between an item of conscious and unconscious knowledge is not like that between the table in front of the curtain and the table behind it.[27]

His use of 'paradigm' here is obviously not equivalent to 'stereotype' but rather to standard case (as opposed to the extended cases). I think this position is quite counter-intuitive. The fact that it takes one person a few moments to remember a given rule (the second stage in Cooper's progression) does not mean that he really didn't know the rule in the same sense as the person who was able to state it immediately. Without any evidence to

the contrary, we would no doubt say that they both knew the rule in precisely the same sense but that the first person couldn't remember what he knew as quickly as the second.[28] If a third person took twice as long as the second to remember, would we say that he 'remembered' in a different sense than the second? Certainly not. And this is because 'remember' is not paradigmatically or criterially tied to such time differences. Similarly, unless we have separate arguments to the contrary, we need not accept Cooper's claim that 'know', 'believe', etc., are paradigmatically or criterially tied to linguistic expressibility or acknowledgement. Cooper, I think, has confused 'paradigm' in the sense in which it is equivalent to 'stereotypic' with another sense of 'paradigm' in which it means the opposite of 'extended'. It would be no shock if we discovered that a cultural historian was able to trace a 'smoother progression' from a Shaker chair to a Chippendale table than from a Shaker chair to the tubing-and-canvass artifact conjured up earlier. Nevertheless, in the important sense, the latter two are both chairs and the Chippendale is most assuredly not a chair. The 'smooth progression' argument I therefore take to be inconclusive.

Taking, however, does not make it so. The feeling might remain that we have sidestepped a real issue. In describing the relation between the speaker and the

grammar as tacit knowledge, we are using the term 'knowledge' (and 'belief') in a new context. I have argued that its use in this new context does not automatically imply that the term is not to be given a new meaning. Nevertheless, we do need some sort of general guideline to explain what contexts are appropriately described as cases of knowledge, tacit or otherwise. If, as I claimed earlier, the limits are not to be set by simply tracing out smooth progressions, then how are they set?

I don't claim to have a comprehensive answer to this question, although the history of philosophy from Plato to Wittgenstein is rife with attempted solutions.[29] Here I think we would do well to return to Freud. The psychoanalytic theory developed by Freud clearly involves a number of paradigm-extensions. The idea of an unconscious mental state or process is but one example and perhaps not the most important or innovative. As Flew points out, the idea that unconscious desires and motivations could <u>cause</u> non-voluntary and non-intentional movement (behavior)--e.g., ticks, impotence, etc.--might be even more radical. The passages cited earlier present more than an argument in favor of unconscious mentation; they hint at a general conception of paradigm-extension: a paradigm for extending the paradigm uses of terms like 'belief', 'motivation', 'desire', and so on. For Freud, the important principle is the explanatory value of such

a posit:

> ...on the other hand, they (conscious mental acts) fall into a demonstrable connection if we interpolate the unconscious acts that we infer.[30]

This point is highlighted in a novel argument that draws the analogy beteen our belief that other beings are conscious and the belief in unconscious mental acts in ourselves. The argument, in short, is that we hold the former view because positing such internal states in others organizes and explains their observable behavior. In ascribing unconscious mental states to ourselves, we are simply applying the same form of justified theoretical inference to the best explanation: ". . . in postulating it (an unconscious)[31], we do not depart a single step from our customary and accepted mode of thinking."[32]

To conclude this discussion, we have seen that to the extent that there is a clear notion of a paradigmatic use and of an extension of such a use, no argument against the tacit knowledge claim is forthcoming. On the contrary, the clearest notion of what should count as knowledge/belief and what should not, rests <u>not</u> on surface similarities to the stereotypical cases, but rather on the explanatory value of such posits. It is for this reason that I think Freud has more than simple historical interest.

The Evidential Connection. We come now to the third conception of the relation between belief-ascription, on the one hand, and avowal, recognition from the inside, etc., on the other. On this view, the latter are not analytically or criterially related to the former; they merely provide evidence for such ascriptions. This conception of the relation is completely compatible with ascription of tacit knowledge of grammatical rules to speakers. If the subject offers p as what he has accepted all along, then we tend to ascribe p to him more or less automatically. But such behavior should not be thought of as a sufficient or a necessary condition. It is not sufficient because, among other things, we must also assume that the avowal or acknowledgement is sincere. It is not necessary, I have been arguing, because ascription can be legitimate without it. Perhaps the special status often accorded to the linguistic manifestation of belief is due to the fact that in day-to-day interaction with people, such expression creates a presumption of belief.[33] But just as we might prove someone guilty even where we have a presumption of innocence (and vice versa), so too we can ascribe knowledge and belief without any linguistic expression of such belief, and we can withhold such ascriptions notwithstanding such expression. Linguistic expression might then be viewed as paradigmatic evidence for belief, but it does not follow

that it is the only sort of evidence.

There is one curious fact that runs through all of this which was noted earlier but bears repeating. The attempt to link ascription of belief to some form of observable behavior is usually part of an anti-mentalist/pro-behaviorist program. In light of this fact, it is strange that so much weight should be given to linguistic behavior at the obvious expense of other forms of evidence. The Quinean picture, which is the most influential version of philosophical behaviorism, would seem to lead to a different conclusion entirely--i.e., that there is no direct correspondence between any one evidential statement and a theory (in this case our own psychological theory of the speaker). Deciding what we take the subject's beliefs to be depends on all the evidence we have about what he does--not just what he does with language. This apparent discrepancy will be taken up in the next section. Before moving onto that, however, I want to make one last foray into the Freudian framework to deal with two very specific questions.

Freud and tacit knowledge: two problems.

We are in a funny position. Our analysis of the Freudian stance leads us to agree with Nagel's conclusion--i.e., that speakers have tacit knowledge of gram-

matical rules--but for reasons incompatible with Nagel's own. Nagel thinks that speakers have the potential for the 'inner recognition' he wants to make a necessary condition for tacit knowledge ascription and that the ascription is therefore legitimate. I have argued that even if speakers have no real potential of this sort (and I don't think they do), tacit knowledge ascription is still legitimate because 'inner recognition' is not a necessary condition for belief-ascription. I have taken the argument a step further into the enemy's camp by adducing evidence to show that for Freud himself, inner recognition is not necessary and was never considered to be the grounds for ascription of unconscious mental states to begin with. I do not claim to have proved anything in respect to Freud's views on this matter, but I think that the passages cited speak for themselves. In effect, then, I have used Freud to support the posit of tacit knowledge of grammatical rules by arguing that such a posit exploits the same form of theoretical inference Freud justifiably adopted.

But now it might be argued that there are two crucial differences between the Freudian and the Chomskyan cases that have nothing to do with the possibilities of recognizing the non-conscious material from the inside. These differences might be considered good reason for accepting the Freudian ascription of unconscious mental

states but rejecting the Chomskyan ascription of tacit knowledge. The first difference is that in the Freudian framework, the fact that the unconscious mental states are unconscious is given a motivational explanation. Freud introduces the notions of resistance and repression, as well as various censoring mechanisms to explain why the contents of these states and processes have not made their way into the conscious mind of the subject. To my knowledge, Chomsky nowhere offers a parallel explanation for the 'tacitness' of our knowledge of grammar. Second, it has been suggested that in the Freudian case, the ascriptions of unconscious states rests on the assumption that the subject in question has the requisite conceptual equipment to formulate such beliefs, fears, and so on. When the therapist ascribes to the patient the unconscious belief that his father killed his mother, the therapist assumes that the patient has the intellectual wherewithal to formulate such a belief. When the psychologist ascribes tacit knowledge of Tag Formation Rule discussed earlier,[34] he does not and, more importantly, cannot make such an assumption. At the very least, the argument goes, we cannot assimilate the Freudian context to the tacit knowledge case.

It is clear that to treat these questions in any detail would involve a very long and complicated detour into Freudian theory and metatheory. Such a detour would

be beyond the scope of this thesis; I therefore limit myself to a few brief comments.

First, a general point: I have not argued that the grammatical context can be assimilated to the analytic context or that Chomsky's notion of tacit knowledge is equivalent to Freud's notion of unconscious knowledge. I have only argued that the same theoretical principle that supports positing the latter works for the former as well. That there are differences is not surprising and not incompatible with my position.

Even so, there are still two points that can be made about the notions of resistance and repression. One is that it is not clear that Freud was content to explain the unconsciousness of all unconscious states in these terms. Consider, for example, the following remarks:

> The truth is that it is not only what is repressed that remains alien to consciousness, but also some of the impulses which dominate our ego and which therefore form the strongest functional antithesis to what is repressed. (Italics mine)[35]

> The content of the Ucs (unconscious) may be compared to a primitive population in the mental kingdom. If inherited mental formations exist in the human being--something analogous to instinct in animals--these constitute the nucleus of the Ucs.[36]

Both these passages (especially the former) indicate that for Freud the unconsciousness of at least some uncon-

scious states has little or nothing to do with repression. Nevertheless, the fact that we have no explanation for the unconsciousness of these other states does not undermine the view that they exist. Instead, it remains a problem that we hope will be solved as we learn more about the mind.[37]

This brings me to the second point. It is quite possible that we will be able to one day explain, in terms of a broad theory of the general organization of the human mind, why knowledge of the rules of the grammar are "beyond the level of actual or even potential consciousness."[38] At least there is no reason to think that such an explanation is out of the question. It is interesting in this regard to consider another view of Freud's (that might be thought of as his Copernican/Newtonian revolution): he suggests that all mental states and processes are by their nature unconscious--what needs to be explained is why some can/do become conscious. [39]

The second question regarding the assumption of intellectual adequacy raises a difficulty that deserves separate treatment. Here it must suffice to point out that although it poses problems for the psycholinguistic ascription of tacit knowledge of grammatical rules, these problems also arise in the psychoanalytic setting--for example, in the Freudean theory of infantile sexuality.

So here again, it is not obvious that we have a significant disanalogy.

Let us return now to the main line of our argument. We recall that a discrepancy was noted between the general aims of philosophical behaviorism--especially the Quinean commitment to a 'holistic' relation between evidence and theory--and the excessive reliance placed on linguistic behavior as a criterion for ascription of belief. We have also found that so far none of the arguments attempting to establish anything more than an evidential relation between belief-expression and belief-ascription is sound. If the connection is no stronger than this--i.e., if it is not analytic or criterial (in some special sense of the term)--then it seems that we have no a priori challenge to the Chomskyan ascription of tacit knowledge. All we are left with is something of a puzzle: how to explain the excessive reliance on linguistic expressions of belief--i.e., avowal, recognition, acknowledgement from the inside, and so on.

The solution to this puzzle will emerge, I think, when we consider one last argument in favor of accepting something close to Nagel's condition on belief-ascription. The argument again tries to show that some sort of linguistic expression (or potential for such expression)

is a necessary condition on belief-ascription. We have already disposed of two attempts to establish this conclusion--i.e., that the connection is (a) analytic and (b) criterial--but this argument approaches the problem from a different perspective. Essentially, the view is that ascription of tacit knowledge of a grammar is unsound on methodological grounds--that it leads to conceptual confusion in theory construction. Let us see why such a claim might seem plausible.

Methodological Scruples: the Problem of Extensional Equivalence

Consider for the sake of argument that we've accepted Chomsky's view of grammar as a theory of competence--i.e., a specification of a system of internalized rules. Imagine now that a linguist has come up with a satisfactory grammar for English; call it G_k. If Chomsky is right, then G_k will have the form of an axiomatized theory. Now if we were convinced by my counterarguments of the previous section to waive Nagel's recognizability requirement, we would hypothesize that the ordinary speaker of English has tacit knowledge of the rules of G_k; no matter that he cannot state these rules or recognize them.

An embarrassment of riches

But now, the argument goes, we are headed for a methodological quagmire. It is a well known fact that if a theory is axiomatizable at all, it is axiomatizable in a number of different ways, not all of which are synonymous or notational variants of the original.[40] All of these axiomatizations are by definition extensionally equivalent. Applied to our own case, since grammars are, according to Chomsky, axiomatic theories, G_k has a corresponding set of extensionally equivalent counterparts $(G_1...G_n)$[41] that are not synonymous with G_k even though they generate the same output.

What we have then, is a metatheoretical guarantee that even if we do discover a grammar of a natural language, we will be saddled with what Stitch has called 'an embarrassment of riches'. How then are we to decide what it is that the speaker tacitly knows? Which member of the set $(G_1...G_n)$ are we to say is internalized? It clearly makes no sense to pick one member of the set arbitrarily and say the speaker tacitly knows that one. The only other possibilities are to ascribe tacit knowledge of all of them, or to look for more evidence that will favor the ascription of one. The argument now will be that both these possibilities are at bottom incoherent; we are therefore forced to abandon the view that speakers have any tacit knowledge at all.

Referential opacity and knowledge ascription. I think we can agree at the outset that the claim that the speaker tacitly knows all extensionally equivalent grammars is inappropriate. One of the central features of normal ascriptions of knowledge and belief is that they are referentially opaque--in other words, they are sensitive to purely intensional distinctions. Proverbially, even if Ralph knows that Ortcutt is the man with my umbrella and 'the man with my umbrella' is extensionally equivalent (in this case: coreferential) to 'the head of the Scottish Secret Service', it does not follow that Ralph knows that Ortcutt is the head of the Scottish Secret Service. If attributions of tacit knowledge did not meet this constraint (which we might think of as an application of Brentano's fourth condition of adequacy)[42], then they would not be knowledge ascriptions in any standard sense.

It might be thought that we have an obvious reply here. After all, not all knowledge ascriptions are referentially opaque. For example, I might tell you that John believes that the guy who hit a grand slam yesterday is the best shortstop in the league. I put it this way because I have simply forgotten the players name. Here the ascription could be true even if John knows nothing of yesterday's game. This might be taken to show that referential opacity is not a characteristic of all belief

ascriptions--and it opens the door for the view that tacit knowledge ascriptions are of this referentially transparent sort. So why can't we say that if A has tacit knowledge of G_k, then he has tacit knowledge of every G_i where G_i is extensionally equivalent to G_k? Extensionally equivalent grammars would now create no difficulties--they simply afford us more ways for characterizing what it is that the speaker knows.

What is wrong with this response is that it mistakes what are really two different sorts of intentions we might have in ascribing a belief for two different types of beliefs. Beliefs themselves are neither opaque nor transparent. A and B might have the same belief, but I intend my ascription of the belief to A to be opaque and my ascription to B to be transparent. I might know more about A than about B and feel that I have more evidence about the exact form of his belief.[43] So even if in ascribing tacit knowledge of a grammar we intend our ascription to be transparent, it makes little sense to say that the mental state--the tacit knowledge of the rules--is itself transparent. If we don't know the exact form of a belief, then we must satisfy ourselves with an attribution that has no more than a referentially transparent force. Be that as it may, the subject must have one belief or the other; as Butler says, "everything is what it is, and not another thing." So it makes no sense

to ascribe all the members of $(G_1...G_n)$ as what the speaker knows tacitly and still contend that tacit knowledge is 'full-blooded' knowledge.[44]

'All the possible evidence'. We now turn to the other possibility. Why can't we say that in fact the speaker tacitly knows and is following one of the grammars in the set $(G_1...G_n)$; it is only that we are empirically unable to determine which one it is? Here is Cooper's response to this:

> ...if we have nothing but RFB[45] (rule following behavior) to go by (and the functionalist admits we do not), then there can be no reason to attribute to the person the belief that R is a rule of his grammar rather than the belief that R1, or R2, ... or Rn is a rule. Any one of these beliefs would have the same consequence as any other. (Italics mine) [46]

His claim is that "our inability to identify the speakers' beliefs or knowledge is not due to any contingent limitations on our powers of investigation" but on the fact that all extensionally equivalent grammars are ex hypothesis equally confirmed by the 'rule following behavior' which is the only possible source of evidence. Employing Quine's principle of 'no entity without identity', Cooper concludes that no states of knowledge or belief can be postulated unless there are identity conditions for them. If there is no way of telling what a person knows or believes, then no knowledge or belief should be attributed to him."[47]

This is a very strange conclusion. Broadly speaking, we have been trying to defend the view that part of the explanation of the speaker's knowledge of language involves ascribing tacit knowledge of a grammar to him. This last argument sets out to show that because of methodological problems, such a strategy is conceptually doomed: we will have no way of choosing between extensionally equivalent formulations. We therefore have to give up the idea that in constructing grammars we are developing theories of what the speaker knows. Instead we must be satisfied to say that all the members of the set ($G_1 \ldots G_n$) are true of the speaker--i.e., characterize his English competence--but not that the speaker has internalized or has tacit knowledge of any one of them. As Quine would have it: although all the systems of rules '<u>fit</u>' the behavior of English speakers, none of them '<u>guides</u>' such behavior.[48]

But aren't we here abandoning our <u>only</u> hope of an explanation? It is as if we had two suspects in a coroner's investigation of a death, and we had sufficient evidence to indict either one of them--e.g. both had a motive, both were seen at the scene of the crime, both left fingerprints on the gun, and so on. If the coroner could not decide whom to hand over for trial, should we conclude that the victim was not murdered at all? It might come to a point where both suspects are reluctantly

released. Still, this is clearly a case of murder; a bullet hole in the heart is certainly not an indication of a natural death. The coroner's verdict therefore ought to be murder (perhaps 'by persons unknown'). What good would it do to adopt the Quinean stategy and say that the evidence fits the hypothesis that the first suspect is guilty as well as the hypothesis that the second suspect is guilty? Analogously, what can possibly be gained if we reject the idea that <u>any</u> member of the set ($G_1 \ldots G_n$) is tacitly known by the speaker? How then will we explain his behavior?

Our opponent will here reply that we have missed the crucial point. Again, the difference between the murder case and the grammar case is first, that in the former instance we have simply run out of evidence that might serve to indict one suspect and exculpate the other. Still, no one could think that on all the possible evidence the conclusion would be the same. The coroner's failure to hand over a suspect for trial is a matter of the paucity of empirical data. In the psycholinguistic case, the problem is purportedly that the candidate hypotheses are all equally confirmed by <u>all</u> the possible evidence.[49] I have already indicated that I don't think the foregoing arguments pose a genuine threat to the program of explaining linguistic functioning by ascribing tacit knowledge. In what remains of this chapter I want

to explain why.

Extensional equivalence and Nagel's condition: the Quinean paradigm. We might start with the following question: Granted that we accept the formal results about alternative axiomatizations and extensional equivalence, what has this got to do with the relation between linguistic expressibility and the ascription of belief? How is it relevant to the claims we have been considering? To clarify this point, I want to look at an influential passage from Quine's "Methodological Reflections on Current Linguistic Theory":

> Imagine two systems of English grammar: one an old-fashioned system that draws heavily on the Latin grammarians, and the other a streamlined formulation due to Jespersen. Imagine that the two systems are extensionally equivalent, in this sense: they determine, recursively, the same infinite set of well-formed English sentences. In Denmark the boys in one school learn English by the one system, and those in another school learn it by the other. In the end the boys all sound alike. Both systems of rules fit the behavior of all the boys, but each system guides the behavior of only half the boys. Both systems fit the behavior also of all us native speakers of English: this is what makes both systems correct. But neither guides us native speakers of English; no rules do, except for some intrusion of inessential schoolwork.
>
> My distinction between fitting and guiding is, you see, the obvious and flat-footed one. Fitting is a matter of true description; guiding is a matter of cause and effect. Behavior fits a rule whenever it conforms to it; whenever the rule truly describes the behavior. But the behavior is not guided by the rule unless the behaver knows the rule and can state it. This behaver observes the rule.[50]

In this case, the fact that we are confronted with a pair of extensionally equivalent grammars leads to no difficulties. One group, we say, was guided by one grammar and the other group by the other grammar.[51] Quine implies that our decision here is facilitated by the fact that the schoolboys can tell us which grammar they use. If they cannot say, at least we have a well-founded expectation that they will remember one when it is presented to them; that they will recognize it as the grammar they learned. Even if these attempts to establish which is which fail, we can always point to evidence of their divergent learning histories: one group was taught Jespersen's grammar; the other was trained from a different Latinlike text. So even though the two grammars are extensionally equivalent, we can point to some real (or at least possible) empirical evidence that will distinguish the hypothesis (J):

(J) 'student s is guided by Jespersen's grammar (G_J)'

from its rival (L):

(L) 'student s is guided by the Latinlike grammar (G_L)'

The point then is that although the grammars they deal with are equivalent in Quine's sense of yielding the same output, the hypotheses are <u>not</u> equivalent--they divide 'all the possible evidence' in different ways.

To put it slightly differently, we can explain the empirical non-equivalence of the hypotheses notwithstanding the equivalence of the grammars by distinguishing the relevant extensions. In calling the grammars G_J and G_L extensionally equivalent, Quine takes the extension of a grammar to be the set of strings it marks as well-formed. Any system which determines this set of strings is extensionally equivalent to G_J and G_L. The 'extension' of a hypothesis, on the other hand, can be roughly conceived of as an ordered pair consisting of the confirmatory evidence and the disconfirming evidence. The extensions of (J) and (L) can now be seen to be different: the fact that student s tells us he was taught Jespersen's grammar and can state the rules is in the positive extension of (J) but in the negative extension of (L).

But now consider again the case of the normal English speaker. Chomsky wants to say that each such speaker has tacit knowledge of the grammar of his language, even if we don't know which grammar to ascribe. For every member of the set of extensionally equivalent grammars ($G_1 \ldots G_n$), we will have a different hypothesis: (H1) The speaker tacitly knows G_1, (H2) The speaker tacitly knows G_2, ...all the way through (Hn). Here then is the crucial difference: in the Danish schoolboy case, we were able to adduce evidence to favor one of the hypotheses

considered by interviewing the subjects and asking them to state the rules. Unless we can exploit the subject's ability to tell us which rules he follows, which grammar he knows, the evidence that determined our original hypothesis--i.e., the data that specified the set of strings an adequate grammar has to generate[52]--is all the evidence he will be able to get. In other words, the extensions of all the hypotheses H1...Hn will be identical; we will never, in principle, be able to specify what the speaker tacitly knows.

Cooper's argument. As we shall see, it is a moot point whether Quine himself wants to draw so radical a conclusion from the foregoing case. For Cooper, though, the lesson is clear:

> Throughout his writings, Quine has insisted on the crucial connection between existence and identity conditions. 'No entity without identity' might be the slogan. Applied to our area of concern, the point will be that no states of knowledge or belief can be postulated unless there are identity conditions for them. If there is no way of telling what a person knows or believes, then no knowledge or belief should be attributed to him.
>
> The 'functional' doctrine, unfortunately, fails to provide identity conditions for the states of belief and knowledge postulated. For, despite claims to the contrary, the behavior which allegedly results from such states cannot serve to identify them, to distinguish one from another. Certainly 'functionalists' are well aware of the importance of identifying the states. Thus Fodor's insistence that
>
> > "the hypothesized psychological constructs are individuated primarily or solely by reference to their alleged

> causal consequences."
> But this could only be true, if at all, if 'behavior' or 'causal consequences' include avowals and expressions of knowledge or belief, or the sort of 'acknowledgement' behavior recently discussed. Behavior and consequences of these kinds, however, are precisely what 'functionalists' cannot have in mind, since they are not to be found among typical, fluent speakers of a language.[53]

Cooper's view is that extensional equivalence becomes a problem when the knowledge in question is held to be tacit. If the subject could express his beliefs linguistically, the clear implication is that we would have a basis for deciding between alternative grammars. On this view, then, something like Nagel's condition is a methodological necessity.

A response to the problem

My own view is that the argument from the existence of extensionally equivalent grammars does not pose any a priori problems of ascription of tacit knowledge of a grammar. I think this is so for two reasons. First, even if the argument were sound, it would prove too much; it would count as an objection against ascription of explicit knowledge and belief as well. So I think that Cooper is confused--even if the speaker's tacit knowledge of a grammar could be somehow made explicit (e.g., satisfied Nagel's condition), the difficulties extensional equivalence supposedly gives rise to would not be eased.

Second, I think that in the final analysis the threatened methodological problems are illusory.

What makes Cohen run? I think we can admit that there is a modicum of intuitive plausibility to the idea that if we could only get the subject to talk our problems would be solved. In our everyday ascriptions of knowledge and belief we often rely on what the subject can tell us to help us decide between hypotheses that are equally confirmed. Imagine the simplest sort of case where we see our neighbor Cohen running toward the train station and we wonder: What makes Cohen run? We develop two hypotheses on the spot:

> (H1) Cohen is running because he believes that he is on time but that the train will leave early.
>
> (H2) Cohen is running because he believes that he is late and that the train will leave on time.

Both hypotheses would explain the behavior in question but we sometimes find it difficult to decide which one is correct. Assume Cohen will not slow down to talk to us. We might try looking for indirect evidence to decide this issue; we might note that the train has never left early before and that Cohen has no obvious reason for thinking it will now. Let us say, though, that no evidence of this sort is available. So we decide to ask Cohen on his way home, and he replies that his watch was running fast

and he really believed he was late. Given the assumption that he is not a compulsive liar (and a few others, including the assumption that he knows English), we take his avowal as establishing (H2). So uncertainty about third-person belief-ascription is, as a matter of fact, usually settled by eliciting a response from the subject. The same is true of ascriptions of fear, hope, desire, etc.

Belief and avowal. What we often forget is that notwithstanding our reliance on the subject's say-so, such avowal is not conclusive; it is not always the last word. Notice that it is still possible to reject Cohen's explicit avowal and continue to ascribe (H1) as a more accurate picture of what he believes. What if we become convinced that Cohen is a compulsive liar or that he has been hypnotized to respond in some particular way to inquisitive neighbors. It turns out then that, strictly speaking, (H2) explains Cohen's total behavior (his morning running and evening avowal) better than (H1) does only if we supplement it with (H2.1):

> (H2.1) Cohen's reply is sincere, he is not hypnotized, he is not a compulsive liar, etc..[54]

But notice also that we can supplement (H1) with (H1.1):

> (H1.1) Cohen is a compulsive liar

and still explain the new total behavior. It will be replied that we can adopt (H1.1) if we want to, but unless we have special evidence we would be silly to do so. The evidence is most simply accommodated if we steer clear of these fanciful assumptions and take Cohen's avowal at face value. This is certainly so. But this does not at all imply that avowal provides incontrovertible evidence for belief-ascription. It could turn out that Cohen is a compulsive liar, and he really did believe, for whatever reason, that the train would leave early.

'Holism' and mental state ascription. The last point is not a new one: ascriptions of mental states, even when bolstered by the subject's explicit avowal, have no claim to apodictic certainty.[55] To adapt the famous Quine-Duhem thesis, we may say that our statements about the internal world face the tribunal of sense experience not individually but only as a corporate body.[56] We can adopt (H2) and (H2.1) and assume that Cohen is like the rest of us or opt for (H1) and (H1.1) and take him to be a compulsive liar. This point about ordinary belief-ascription can be applied to the grammatical case as well. The implication of Quine's example was that the potential avowal of the Danish schoolboys would settle between hypothesis (J)--that student s is guided by G_J--and hypothesis (L)--that student s is guided by G_L.

But let's say that the student can state the rules of Jespersen's G_J. If we want to be stubborn, we can defend a comprehensive hypothesis that includes the far-fetched assumption that he really uses G_L but has been taught to lie to foreigners. In point of fact, doubting the veracity of a randomly selected Danish schoolboy is like assuming that Cohen is a compulsive liar--both are logically impeccable but practically ludicrous. But to show that an extensionally equivalent hypothesis is ludicrous is not to do away with the methodological problem. After all, we have no metatheoretical guarantee that every grammar will have non-ludicrous extensionally equivalent counterparts; no guarantee that we will be as comfortable with each of the counterparts as with the original.

A reply to Cooper. So, even if we are given all the data involved and all the details of the subject's avowal-behavior, there will still be extensionally equivalent hypotheses which explain all the evidence. We would only be able to use avowal (recognition, etc.) to solve our problems if we construe it as somehow standing apart from the other evidence that can be relevant to belief-ascription. If we develop alternative hypotheses that account for all the data except the subject's linguistic behavior, and we then introduce avowal, then we might rule out a number of alternatives as incompatible with the new evidence. But this is only a complicated way of

saying that the more evidence you have, the better are your chances of eliminating bad hypotheses. Of course if you put some of the relevant data in a "Don't open till all the other evidence is in" envelope, then it will become especially important. But why should any evidence be so honored? Cooper thinks that avowal is the great remedy for the problem of extensional equivalence only because of the unnatural division he makes between avowal-type evidence and the evidence having nothing to do with avowal. One could, it seems, make the same sort of argument to show that evidence culled from entries in a personal diary deserve special status because they too can help us decide between otherwise equivalent hypotheses. The possibility of avowal only postpones the issue of extensional equivalence; it does not in any way resolve it. Therefore, even in cases where this possibility can be presumed to be absent--e.g., tacit knowledge of a grammar--it is not enough to point out that there are extensionally equivalent hypotheses to be had. These are always available.

If we reject Cooper's attempt to establish a special place for the speaker's avowal and accept the foregoing characterization of the relationship between evidence and belief-ascription--a characterization that, ironically, owes much to Quine--then I think we must agree that Cooper's argument for requiring belief and knowledge

ascriptions to meet Nagel's recognition condition fails.

So far I have only dealt with half the challenge. I have argued that if extensional equivalence does in fact pose a methodological problem, then that problem has nothing to do with tacit knowledge as opposed to knowledge that is explicit. We might say that if tacit knowledge ascriptions are subject to the problem of extensional equivalence, it is not because of their 'tacitness' but because of their intensionality; that is, because psychological explanations in terms of ascriptions of mental states like belief are intensional.[57] It goes without saying, of course, that an explanation is intensional irrespective of whether or not explicit formulations of the contents of the ascribed states are potentially recognizable by the subjects in question. This is why I said earlier that the argument really justifies the conclusion that if we can't decide conclusively between extensionally equivalent hypothesis-grammars then we ought not to say any one of them is tacitly known, then it would also justify the conclusion that we ought not to say anyone has any explicit knowledge (belief, etc.). So it is not only Chomsky's position that must come to grips with 'the problem' of alternative equivalent hypotheses; any position that explains behavior in terms of underlying intensional states--this includes common-sense explanations and mentalistic cognitive psychology--will be

in the same situation.

But someone can accept my argument and try to turn it on its head. In other words, if extensional equivalence presents a challenge to the explanatory power of all intensional ascriptions, then it surely challenges the explanatory power of tacit knowledge ascriptions. It is no defense against this sort of bullet-biting to say that ascriptions of tacit knowledge of a grammar are only as methodologically incoherent as is, say, cognitive psychology.

For one familliar with Quine's behaviorist views on psychology and reductionist views on psychological explanation, it is natural to think that by 'someone' I mean Quine. This broad skepticism with regard to all intensional explanations seems to be wholly in character, so for the time being we will take this to be Quine's position.

Meeting the second half of the challenge I spoke of earlier would involve showing that this sort of broad skepticism is unjustified. Let me say at the outset that I will not take on this task here; to do it justice would require a general defense of cognitive psychology, and such a defense is beyond the scope and requirements of the present work. What I will do here is show first, that this sort of skepticism cannot be at the heart of

the Quinean objection to tacit knowledge ascriptions, and second, that Quine's real objection is based on unrealistic assumptions about the range of evidence relevant to grammars.[58]

Quine and tacit knowledge: three possible perspectives

The crucial consideration I want to appeal to regarding the first point is that Quine commits himself to some sort of notion of rule-following in his discussion of the case of the Danish schoolboy. What I will try to show is that such a notion is neither more nor less problematic than the ascription of tacit knowledge of a grammar.

Skepticism with regard to intensional explanations. Let us assume then, for the sake of the argument, that Quine is skeptical about all intensional ascriptions. Now consider again the case of the Danish schoolboys. Quine's position there seemed to be that where we have explicit avowal we can say that one or the other grammar (but not both) guides the behavior of any given schoolboy. But rule-guidance of this sort seems to constitute a clear case of intensional explanation. Quine says explicitly that "each system guides the behavior of only half the boys," even though all the boys have knowledge of the language that conforms to both systems. If the

real objection to tacit knowledge ascriptions is that they are intensional, then how is such talk of rule-guided behavior to be accommodated? How can he ascribe such rules in the face of the possibility of extensionally equivalent alternative hypotheses to explain the behavior?

Intensionality and rule-guidance. It might be replied that this problem is the result of our reading too much into Quine's claim that the behavior in question is 'rule-guided'. Perhaps explanations in terms of guiding rules are not intensional, after all. Perhaps all he means by 'guiding' is that the student's behavior fits the rule and the student can state the rule. On this interpretation, Quine is not ascribing an internal system of rules that plays a role in the explanation of the schoolboy's linguistic knowledge; this would indeed leave him open to the problem of extensional equivalence. Instead, the point would be that we can, if we like, provide a behavioristically acceptable replacement for the mentalistically tainted notion of rule-guidance--i.e., conforming behavior plus avowal.

Closer attention to what Quine says a propos of his example will convince us that this is not his intent. Recall that in the ensuing discussion of the case, he provides two different characterizations of rule-guid-

ance:

> (1) "guiding is a matter of cause and effect"
>
> (2) "the behavior is not guided by the rule unless the behaver knows the rule and can state it"[59]

Although Quine does not point this out, the two characterizations are <u>not</u> equivalent--at least not patently so. Perhaps on (2) we might make the case that Quine is simply providing the best behaviorist counterpart for the mentalist notion; (1) puts the lie to this suggestion. It seems that at least where there is the real possibility of avowal, Quine is willing to explain some of the subsequent behavior of the subject in terms of the intensional notion of 'being guided by a rule'. So we still have our incompatibility. How are we to square this causal picture of rule-guiding--which sounds very much like a cognitivist explanation--with the earlier skepticism with regard to intensional explanations?

Perhaps they an be fit together in this way. When Quine says that there is a causal sense in which the rules guide the speaker, he may still not mean that the subject has an internalized system of rules--an 'on-board' computer--that is in any sense 'Jespersenic' as opposed to 'latinate'. Although this latter sort of position <u>would</u> be an instance of theory construction along mentalist/intensionalist lines, Quine need not accept it.

What Quine might really mean is that to explain the speaker's ability, we could give a causal history of its genesis. To do this we have to include the learning experience of the student, and only here will we have to mention one particular grammar--G_J or G_L. The student's learning history is 'Jespersenic' or 'latinate'; what he has now (a disposition) is neither. The only way one member of a set of extensionally equivalent grammars 'guides' behavior is if a token of that grammar was causally implicated in the acquisition of the linguistic skill.[60] It seems, then, that Quine can escape our inconsistency charge after all. His view would be that rule- guidance is a matter of explicit training. Where we have evidence that the training has occurred, we can say that rules played a causal role. In the case of the purported tacit knowledge Chomsky wants to ascribe to native speakes, everyone is willing to grant that there is no (evidence of any) explicit instruction in any grammar. The notion of 'guidance' is therefore out of place here; it can be given no empirical content.

We are still left with a question: How does citing details of the instructional history of the pupil explain anything? How does the training in the past cause the present behavior? Only the most radical Skinnerians, and in this regard Quine is to be excluded, would accept this sort of 'action at a distance' explanation as final. The

Quinean perspective is that the training establishes certain dispositions which are ultimately to be cashed in in terms of physical structures.[61] The structures, then, play a mediating causal role in the production of verbal behavior and are distinct for different grammars G_J and G_L. So perhaps we should think of rule-guidance as a physically defined phenomena. A subject is taught a set of rules R and the result of this instruction is that the subject comes to be in physical state S. Physical state S will figure in the causal explanation of the behavior B(R) that is enjoined by the rules R. If the subject had been taught an extensionally equivalent set of rules R', then one can assume that this would have eventuated in a different physical state S' which again would have had its causal consequences B(R'). The sets R and R' are extensionally equivalent for Quine because B(R)=B(R').[62] However, $\sim$(R=R') and $\sim$(S=S').[63] So it seems that we can explain rule-guidedness in terms of behavior, tokens of rules, and physical states.

But now I think Quine is again in trouble. This stance concedes <u>exactly</u> what we assumed he wants to deny. If we conceive of rule-guiding as neurophysiologically determinate, then we will have to admit that it's possible that just these neurophysiological structures are present in <u>native</u> speakers that were never explicitly trained to use <u>any</u> specific grammar. How could Quine

then reasonably deny the claim that a native speaker's linguistic behavior is guided by rules of a grammar?[64]

Response to Quine: a neurological fantasy. To bring home this point, consider the following supplementation of the Danish schoolboy example. We have two sets of students--call them 'S_J' and 'S_L'. Both sides in our dispute agree that as far as these subjects are concerned it makes sense to talk of their subsequent behavior as rule-guided. Now suppose that a new student arrives who is a native speaker of English. Chomsky wants to claim that this student, unlike his schoolmates, knows a grammar of English tacitly. According to Chomsky's critics, this student knows no grammar at all; he is not guided by any rules even though his linguistic behavior fits G_J and G_L equally well. Now imagine that a student taught G_J, a student taught G_L, and the new student all die, and we perform autopsies on the three brains. Much to our surprise, we find that the new student's brain bears a remarkable resemblance to the brain of the G_J-trained student in just one area,[65] but differs radically from the G_L-trained student's brain in just that spot.

This is very meager evidence, but we can enrich it. Let us assume that we have taught one thousand students English using G_J and another thousand using G_L. We find on autopsy that the brains of all one thousand members of

the G_J group are similar to each other and to the brain of the new student; the brains of the G_L group also exhibit the same uniformity but are markedly different from the brain of our new student. It seems to me that we have some empirical evidence that favors the hypothesis that the new student had knowledge of G_J but that he was unable to discover this by introspection--i.e., he only knew G_J tacitly. I am not claiming that this ascription has the status of a certainty. It is an empirical claim that could be disconfirmed--perhaps all the members of the G_J-group and the new student took a certain form of LSD that the students studying the old fashioned G_L never tried. Be that as it may, the story as it stands, if true, would provide empirical support for the claim that the native has tacit knowledge of one grammar (G_J) but not of an extensionally equivalent counterpart.

To be fair, itmust be admitted that the example is so far-fetched as to admit an indefinite number of alternative explanations for the imagined results. One rather questionable assumption made is that there is very little difference between the acquisition of a first language by normal exposure and the acquisition of a second language by explicit instruction; or at least that similar neural structures would be implicated in both processes. I have no reason to think this is so. I

have also assumed, since I am arguing against a Quinean position, that some form of the psychophysical correspondence theory is accepted, and that it is not a functionalist identity theory. The assumption is not crucial and I have my doubts about its truth. What I want to stress, however, is that I am not arguing that neurophysiological theory will one day settle all our questions about tacit knowledge and psycholinguistics. As I indicated earlier, even if we have 'all the neurophysiological evidence there is'[66] we will still have to make an inductive leap in ascribing tacit knowledge of a specific grammar. I only claim that if these hypothetical findings would count as evidence, then the question of which of a set of extensionally equivalent grammars we say the speaker tacitly knows would be an empirical question. In other words, the question cannot be ruled out of court on the basis of a priori methodological considerations. I have also argued that Quine's (and Cooper's) conception of rule-guidance as a causal notion commits him to the claim that neurophysiology is relevant. So, I conclude that the argument from extensional equivalence fails.

Summary. Let me summarize this rather convoluted dialectic. We began by arguing that if explanations involving tacit knowledge ascriptions were conceptually problematic, then so were all explanations in terms of underlying mental states, the reason being that the

problem of extensional equivalence applies to ascriptions of tacit knowledge because of their intensionality--not their tacitness. We then considered the possibility that Quine would wholeheartedly accept this conclusion and reject all explanations involving ascriptions of intensional states. Such a step, however, seemed incompatible with his willingness to grant that, in cases where there is a possibility of explicit avowal, we can talk about genuine rule-guided behavior--some intensional explanations do seem methodologically acceptable on Quine's view. It was conjectured that perhaps all Quine means in saying that a set of rules guides behavior is only that learning the rules played a causal role in the development of the brain structures that mediate the linguistic behavior in question. But to this it was countered that the same brain structures could be present in a native speaker who has never learned the rules explicitly; at least there is no a priori reason to think that this is impossible. If such evidence were possible--and the autopsy story was just one such scenario--then Quine would have to admit that even without avowal the question of whether a grammar plays a causal role in the etiology of behavior is empirical. We concluded, then, that problems about extensional equivalence pose no special a priori difficulties for ascription of tacit knowledge of a grammar.

Methodological reflections on tacit knowledge ascriptions. The same conclusion might be arrived at from a slightly different perspective. We could say that Quine, unlike Cooper, does not raise the problem of extensional equivalence as the basis of a knock-down a priori argument against the ascription of tacit knowledge. Quine's point is simply that ascriptions of knowledge must rest on clear behavioral criteria. In this regard, he is not systematically opposed to explanations ascribing knowledge and belief to a subject even granting that such ascriptions are intensional. Such ascriptions generally pose no problems because even where p is extensionally equivalent to q we can distinguish the belief that p from the belief that q behaviorally: in terms of the disposition to verbal behavior. Such dispositions are admittedly rough and ready and are only more or less trustworthy. Still, they do constitute an empirical criterion. Such behavior--avowal, acknowledgement, etc.--does not give us a direct reading of the subject's beliefs. What it does give us is a readily applicable criterion for making at least a principled decision between alternative ascriptions. As I stressed earlier, there is no guarantee that we won't be left with a set of extensionally equivalent hypotheses to explain all the data; quite the reverse. Extensionally equivalent alternatives will always be with us.[67] What the

verbal dispositions do is give some empirical sense to the claim that the subject knows q but not p even though they are equivalent.

Quine's misgivings. What puzzles Quine about tacit knowledge claims is that there are no verbal dispositions to fall back on. Consequently, there seems to be no empirical sense to the claim that the subject believes p but not q or vice versa. What Quine wants here is not a philosophical guarantee that we will hit upon '<u>the</u>' right hypothesis; he is simply asking for a methodological policy-decision that makes sense. If we have a set of extensionally equivalent candidate grammars and we have no avowal to go on, how is the linguist supposed to decide which to adopt? On this construal, Quine is literally offering only methodological reflections, not arguments; to wit,

> Implicit guidance is a moot enough idea to demand some explicit methodology.[68]

If we take this perspective on Quine, then what we really need is not a counterargument but assurance. The autopsy story can be taken to show that <u>apart</u> from the dispositions to avow, to acknowledge, etc., there <u>are</u> other ways--for instance, neurophysiological--to distinguish between extensionally equivalent grammars. In other words, we can still give empirical sense to the claim that a speaker has knowledge of G_K but not of the exten-

sionally equivalent G_M even though the knowledge is tacit.

From either perspective on Quine, the conclusion will be the same. Quine's problem is that his view regarding the evidence available for judging the adequacy of grammars is unrealistic. The only evidence he ever mentions as relevant is the set of utterances the linguist observes and whatever avowal the speaker makes.[69] In the Danish schoolboy case, the first set of evidence would have helped narrow the possible hypothesis-grammars, the second would have given a solid reason for ascribing (say) G_J. In the native speaker case only the former is available. But again, why should only these sources count as relevant?

Chomsky's assurances. As Chomsky has pointed out, the range of possible evidence is not even limited to neurological findings of the sort imagined earlier. In general, we can distinguish levels of empirical adequacy that hypothesized grammars can be required to meet. To take one prominent example: Quine, as we've seen, only considers one form of equivalence--i.e., extensional. (Again, two grammars are extensionally equivalent if and only if they determine the same set of well-formed sentences.) Now we might have two grammars equivalent in this sense that can still be empirically distinguished in

another sense. For instance, the first, (G_1), might provide one structural description for a string S while the second (G_2) provided a different one. Psycholinguists have given us good reason to think that there will be detectable empirical differences between a speaker who tacitly knows G_1 and a speaker who tacitly knows G_2. In a celebrated experiment, subjects heard sentences through one earphone, and clicks through the other; in reporting at which points in the sentences they heard the clicks, subjects tended to displace the clicks to the 'phrase boundaries'.[70] Now imagine that the assignment of structural descriptions of G_1 fits the subjective phrase boundaries of the click experiment while the ascription of G_2 does not. Why shouldn't this count as empirical evidence for the hypothesis that the speakers are using G_1 rather than G_2? This evidence in itself might not be sufficient to warrant ascription of G_1 as what the speaker tacitly know, but why deny that it is evidence?

It might be the case that there is more than one extensionally equivalent grammar that can accommodate the click data. Nevertheless, once it is admitted that click phenomena (or the hypothetical neurophysiological evidence) are relevant, the experimental floodgates are opened. As I hinted earlier, there is no real theoretical limit to the evidence that could be pertinent to an

empirical hypothesis. For any set of extensionally equivalent grammars (in Quine's sense) that we think might be internalized in the speaker, we can ask at least the following questions.

Do the grammars under consideration:

(1) determine the same set of strings as well-formed?

(2) determine the same set of structural descriptions for strings?

(3) explain speakers' intuitions equally well?

(4) accord equally well with findings of psycholinguists regarding (say) response time?[71]

(5) fit equally well with our best theory of the acquisition of language and the development of a learner's knowledge of language?

(6) fit our best theory of mind in general?

I see no reason why this list cannot be extended indefinitely.

This last point, then, is that even if we adopt a Quinean attitude toward psychological explanation and press for clear-cut behavioral criteria for intensional state ascriptions we can still make empirical sense of the claim that a speaker has tacit knowledge of only one of a set of extensionally equivalent grammars. All we need to do is look to a broader range of empirical phe-

nomena than mere sentence production, and we will be able to make the necessary distinctions. How, we might wonder, could Quine have missed this point? Isn't obvious that if we are developing a psychological theory of the speaker we have to take more into account than the set of sentences he might utter in his daily speech?

A new look for Quine? The answer here, I think, is that Quine has not missed the point at all. Rather, he wants to distinguish the following two questions (which Chomsky takes to be equivalent):

(1) Which set of rules do speakers know tacitly?

(2) What is the correct grammar of a language?

As far as (1) goes, Quine has in the past been a little skeptical about the possibility of answering this question. But the skepticism, although on the wane, is overshadowed by disinterest. For Quine, a language is constituted by a specific set of dispositions to verbal behavior. An individual is a member of a linguistic community, given the appropriate idealizations, if he shares these dispositions with his neighbors. A grammar is, at the very least, a theory of the language--i.e., a theory about these shared dispositions. Although the speaker's mental/brain structures will ultimately account for the dispositions he has, they are not to be identified with the language. What the speaker has 'in his head' is therefore of no special interest as far as

grammar construction and evaluation are concerned. Perhaps we can make the point this way. Let's say we have a speech community--take ours--and we find that each speaker has a different but extensionally equivalent grammar 'in his head'. There is very little likelihood that this will be so given the uniformity of the species and the assumed uniformity of acquisition mechanisms. Nevertheless, the hypothesis is not logically impossible, and it will be interesting to consider it for a moment. For Quine, given that the language continues to play the same communicative role, the fact that we have different 'machinery' inside is irrelevant.[72] As long as the constitutive dispositions remain the same, we are dealing with the same language.

This quasi-functionalist approach might explain Quine's skepticism about the idea of a 'correct grammar'. Since he rejects the Chomskyan psychologization of grammar discussed in Chapter I, a grammar for him will <u>not</u> be a theory of underlying competence; it is instead a theory of the language viewed as a set of dispositions. As a theory it can serve a number of different functions: it might help Danish schoolboys learn English, it might help a missionary contrive a translation manual to make prayer services possible in a native tongue, and so on. It could turn out that one grammar serves the first purpose well but is useless for the second. And here Quine

would say that we can use another of the extensionally equivalent grammars for the second undertaking. In fact, if we are psychologists with the goal of constructing a theory of the speaker's underlying competence, then we might do well to attend to the grammar that the Chomskyan research program finally fixes on. What Quine is arguing against, then, is the appropriation of 'grammar' to cover only this last sort of psychological theory.[73] As he sees the matter, what is in the head of the speaker is no more the grammar than what is in the missionary's notebook. Since grammars (like all else) are practical devices, it makes no sense, for Quine, to disparage any one that accommodates the subject matter--i.e., that determines the same set of well-formed strings.[74] This perhaps is what Quine has in mind in arguing that there is "no truth of the matter" and that there is "nothing to scrute" in the native mind: not that there is no underlying competence at all, but only that the exact nature of this underlying competence is not an essential feature of the language. The latter can vary (over the range of extensional equivalence) and the language--the object of study for linguistics--will remain the same. A theory of language, then, cannot simply be a theory of underlying competence that Chomsky strives for; for Quine this is psycholinguistics but not linguistics.

As I have laid things out, Quine is in accord with the practice of ascribing tacit knowledge of a grammar to a speaker; he is simply unwilling to allow this tacitly known grammar to take on the status of <u>the</u> grammar. One of the caveats I entered in Chapter I was that I would not argue for any classification of grammar construction as a branch of cognitive psychology.[75] My position was simply that an explanation of our knowledge of language must involve reference to our tacit knowledge of linguistic rules. Although I see Quine as agreeing with this last claim, I am well aware that this is an unorthodox interpretation; it might strike the reader as an amalgam of a few Quinean positions with a host of distinctly non-Quinean considerations thrown in. I will grant that there are a number of passages in Quine's work that are incompatible with this reading of his stance on the issues.[76] Nevertheless, I think that Quine has at least been moving toward this position, if he has not already embraced it. To support this contention I offer the following rather lengthy excerpt from a recent paper of Quine's for consideration:

> There is room for confirmation here of the hypothesis that part of the linguist's grammar has structural counterparts to the neural hookup of the speaker, not by examining the nerve net but by studying the infant's progress in the language. We can discern his stages of competence, by noting the periods at which he overcomes various systematic errors and at which he first produces various constructions or responds appropriately to them. Such deve-

lopmental studies should enable us to discover in a tentative way the childs own implicit system of grammatical rules, these being separable from one another by their different times of acquisition.

A lay native speaker cannot state his grammatical rules. He just behaves in the way that the rules describe. By this token any extensionally equivalent system of rules could be ascribed to him indifferently. Chomsky has denied this indifference, postulating rather one or more innate skeletal grammars. I cannot welcome this postulate without further light on the nature of the supposed equipment. We do see a persuasive basis on which the indifference might be challenged however, when we think of a native speaker's linguistic competence as a system of habits that he successively acquires. One set of grammatical rules certainly has an important claim to distinction over other extensionally equivalent sets, if its several rules correspond to temporally successive increments in the native's actual learning of language.

Bever's click technique is an ingenious way of finding out something about one's tacit rules without studying one's development. ...Bever's click experiment is looked upon as confirming the linguist's explicit codification of grammar, by showing that the subject's tacit segmentations agree with that codification...

If a linguist were to compile a grammar for an exotic and hitherto unknown language, it would be interesting to try some clicks and see whether the compiled grammar matched the native's tacit grammar or was merely extensionally equivalent. In the latter event the linguist might care to switch over. I would suggest again that the multiplicity of extensionally equivalent grammars could in theory be put to multiple uses: one grammar might best serve translation into one language, another grammar might best serve translation into another language, yet another grammar might best reflect the psychology of the native speaker, and, conceivably but implausibly, yet another grammar might be simplest. And finally, out beyond extensional equivalence, still further variants could have their uses. (Italics mine)[77]

My remarks here are not intended to show that Chomsky and Quine are in complete harmony. Apart from the differences in approach that can be inferred from this passage, there are other divergencies that have not been touched on here. Among these is the dispute over the appropriateness of mentalistic theories of mind as intermediate explanations of cognitive functioning,[78] and the question of whether language-specific innate mechanisms must be postulated to explain language acquisition.[79] What I have tried to show is that Quine's position on the limited question of ascribing tacit knowledge of a grammar to a speaker is at worst noncommittal, and at best supportive.

We have tagged a number of different objections to tacit knowledge ascriptions as 'Quinean' only to arrive at the view that Quine may not be opposed to such ascriptions after all.[80] To review these objections as they arose: we first dealt with the argument to the effect that <u>implicit</u> rule guidance was an incoherent notion; we then went on to the view that any explanation in terms of ascriptions of intensional states--like tacit knowledge of a grammar--that were not provided with clear behavioral criteria of application were problematic. Finally we reached the contention that what is wrong with Chomskyan ascriptions of tacit knowledge is the fact that they are exclusionary: that they assume that whatever it

is that the speaker tacitly knows is *de facto* the grammar of the language. I don't propose to take sides on this last issue here because it only relates tangentially to the position I set out to defend. My aim in this chapter was to show that there is no sound a priori objection to the claim that speakers have tacit knowledge of linguistic rules. This I think I have done.

NOTES CHAPTER II

1. Rulon Wells ascribes a position something like this one to Locke; I cannot find any evidence in the Essay that Locke held such a view. See "Innate Knowledge" in Language and Philosophy: A Symposium, ed. Sidney Hook (NY: NYU Press, 1969), pp. 101-103.

2. Stephen Stitch presents a similar case in his "What Every Speaker Knows" in Philosophical Review, LXXX, 4, October 1971, pp. 491-492.

3. For an interesting discussion of these issues see Arthur Collins' "Unconscious Belief" in Journal of Philosophy, LXVI, 20, October 16, 1969, pp. 667-680.

4. Thomas Nagel, "Linguistics and Epistemology" in Hook, Language and Philosophy, pp. 175-176.

5. For the purposes of this discussion I will not distinguish between "A tacitly knows rule R" and "A is implicitly guided by rule R".

6. Ibid., p. 176.

7. Ibid.

8. Ibid., p. 175.

9. This rule is taken from Adrian Akmajian and Frank Heny, An Introduction to the Principles of Transformational Syntax (Cambridge, Mass.: The MIT Press, 1975), pp. 221-222.

10. David E. Cooper, Knowledge of Language (NY: Humanities Press Inc., 1975), p.68. The same point is also made by Stitch in "What Every Speaker Knows", pp. 489-590.

11. There is obviously more to Freud's views on the subjects treated here than I will touch on. The essays I discuss are "The Unconscious" (1915) and "A Note on the Unconscious in Psychoanalysis" (1912). Both are included in Sigmund Freud, Collected Papers IV, translated by Joan Riviere (London: The Hogarth Press and the Institute of Psycho-Analysis, 1948), pp. 22-30, 98-137.

12. Ibid., p. 22.

13. Ibid., p. 100.

14. Ibid., p. 99.

15. See also "...the presence of such an idea (unconscious) in the mind admits of indirect proofs of the most cogent kind, which are equivalent to the direct proof furnished by consciousness." (italics mine) Ibid., p. 25.

16. Freud, Introductory Lectures on Psychoanalysis, (London: Allen and Unwin, 1929), 2nd edition, p. 217.

17. Anthony Flew, "Motives and the Unconscious" in Minnesota Studies in the Philosophy of Science, v.1: The Foundations of Science and the Concepts of Psychology and Psychoanalysis, ed. Herbert Feigl and Michael Scriven (Minneapolis: University of Minnesota Press, 1956), p. 156.

18. Nelson Goodman, "The Epistemological Argument" in The Philosophy of Language, ed. John Searle (Oxford: Oxford University Press, 1971), pp. 143-144.

19. See Rudolf Carnap, Meaning and Necessity (Chicago: The University of Chicago Press, 1947), sections 13-15, pp. 53-64, and the discussion in I. Scheffler, Conditions of Knowledge: An Introduction to Epistemology and Education (Chicago: Scott, Foresman and Company, 1965), pp. 76ff.

20. How totally is a matter of divergent opinion, and serves to delineate various points on the logical behaviorist spectrum.

21. For an interesting treatment of these matters see D.M. Armstrong, Belief Truth and Knowledge (Cambridge: Cambridge University Press, 1973), pp. 21-37.

22. Donald Davidson has attempted to provide an argument in support of this view (or at least in support of a closely related view). Unfortunately, I am unable to determine exactly what the argument is. See his "Thought and Talk" in Mind and Language: The Wolfson College Lectures 1974, ed. Samuel Guttenplan (Oxford: Clarendon Press, 1975), pp. 7-23.

23. Cooper, Knowledge of Language, pp. 70-71.

24. Wittgenstein's notion of family resemblance might be considered a non-theory.

25. The context of Cooper's discussion makes it clear that by 'acknowledgement' he intends to cover Nagel's recognition 'from the inside'.

26. I don't want to get into a discussion of what makes

it a chair, although the functionalist tone of my remarks is meant to have broader application.

27. Cooper, Knowledge of Language, p. 72.

28. Recall the game show contestant of Chapter I.

29. Brentano's conditions of adequacy taken up in Chapter I attempt to deal with the same problem.

30. Freud, Collected Papers IV, p. 99.

31. Freud vacillates between taking 'the unconscious' ('Ucs') as an abbreviation or shorthand for the set of unconscious mental conceptions, and viewing it as referring to a mental 'place' or topographical area. See Ibid., pp. 104-109.

32. Ibid., p. 101.

33. See pp. 41ff.

34. See p. 10.

35. Freud, Collected Papers IV, p. 124.

36. Ibid., p. 127. Notice that Freud offers an open-ended suggestion that there are inherited mental formations that are located in the Ucs. Ironically, Chomsky's suggested innate universal grammar can be construed as a partial specification of just such an inherited mental structure.

37. It should be pointed out however, that in the earlier 1912 paper cited above, Freud held the view that the explanatory power of the concept of resistance was broader: "Unconsciousness is a regular and inevitable phase in the processes constituting our mental activity; every mental act begins as an unconscious one, and it may either remain so or go on developing into consciousness, according as it meets with resistance or not." Ibid., p. 27. The implication here is that all unconscious states are unconscious because of resistance.

38. I take it that a why-explanation here will be a specification of the evolutionary value of such an organizational structure.

39. "In psychoanalysis there is no choice for us but to declare mental processes to be in themselves unconscious, and to compare the perception of them by consciousness with the perception of the outside world through the sense organs...." Freud, Collected Papers IV, p. 104.

For a related view applied to cognitive processes (but not to the output of such processes) see Daniel Dennet, Brainstorms: Philosophical Essays on Mind and Psychology (Bradford Books, 1978).

40. Both synonymy and notational variance are notoriously difficult notions. Most discussions of extensional equivalence take it for granted that a clear conception of a 'non-trivial alternative' can be formulated. For the sake of argument I accept this assumption.

41. 'n' here need not be finite.

42. See p. 18.

43. See Armstrong, Belief Truth and Knowledge, pp. 25ff. Cf. Davidson, "Thought and Talk", pp. 7ff.

44. Cooper makes this argument in Knowledge of Language, p. 77.

45. Cooper does not want to admit that this description has a reference; it's more a matter of reference under a false description.

46. Cooper, Knowledge of Language, p. 73.

47. Ibid. See also his "The Deletion Argument" in Midwest Studies in Philosophy, v.2: Studies in the Philosophy of Language, eds. P. French, T. Uehling, Jr., and H. Wettstein (Morrris: The University of Minnesota, 1977), pp. 138-144.

48. W.V. Quine, "Methodological Reflections on Current Linguistic Theory" in On Noam Chomsky: Critical Essays, ed. Gilbert Harman, (NY: Anchor Books, 1974), pp. 104-117.

49. Notice that in the murder case, the hypotheses are not extensionally equivalent--they are simply equally well-confirmed.

50. Quine, "Methodological Reflections", p. 104.

51. Notice that this passage may contain two qualifications that Nagel later ignores. First, he claims that it is "usually possible" to bring the patient to recognize from the inside. What then are we to say of those cases where it is not possible? Can we say that ascriptions of unconscious beliefs and motives ought to cease? Presumably not. I would assume that Nagel's position is that once we accept a context as appropriate for tacit knowledge ascriptions, then we have an umbrella explanatory

apparatus that can be applied freely to all the cases in the context. Secondly, he says that in the psychoanalytic cases "what was unconscious can be brought, at least partly, to consciousness." What of those beliefs that are not so elicitable? Again, I will assume that for Nagel unconscious beliefs need not meet his condition of inner recognition individually.

52. To be precise, the data does not specify any set of strings--the idealization process has a large part to play here. See Chapter II, note 69.

53. Cooper, Knowledge of Language, pp. 72-73.

54. There is some question as to whether or not this 'etc.' can be cashed in.

55. Whether this is also true for first-person ascriptions and for all mental states are complex issues that I will not deal with here.

56. Quine's original statement is found in "Two Dogmas of Empiricism" to be found in his From A Logical Point of View: Logico-Philosophical Essays (Cambridge, Mass.: Harvard University Press, 1953), p. 41. I of course don't mean to imply that Quine would welcome this adaptation.

57. I say 'like belief' because some explanations involving ascriptions of mental states may not be intensional--i.e., when we ascribe non-intentional states (like pain states) to explain behavior. Here it is easy to confuse intensional with intentional.

58. One warning: throughout this section of the argument I will be talking as if Quine is straightforwardly opposed to ascriptions of tacit knowledge of a grammar. Later on I will explore alternative construals of Quine's position, but now I want to take up the 'orthodox-Quine' as he has been read by Cooper, Stitch, Root, and others.

59. See pp. 37-38.

60. Which member of a set of extensionally equivalent rule-systems plays the guiding role then becomes a historical question.

61. For a recent, clear statement of Quine's position on this matter see his "Reply to Lycan and Pappas", Philosophia, 7, 3-4, July 1978, pp. 637-638.

62. 'a=b' should be read as 'a is identical to b'.

63. The non-identity of the two physical states that eventuate from the different trainings is something I assume stands to reason. We might take as evidence the fact that the students in this case would be expected to offer different verbal responses to queries about their training; such differences in behavioral dispositions must, for Quine, signal neural differences.

64. See pp. 37-38.

65. 'area' and 'spot' here are meant to be no more than figures of speech. The hypothetical resemblance I allude to should be thought of in terms of whatever is the best theory of neural similarity.

66. There is some question as to whether this notion makes any sense; more generally, whether we can ever talk of all the possible evidence for a hypothesis. For more on this issue see Chomsky's Reflections on Language (NY: Pantheon Books, 1975), ch. 4. The reader will notice that my discussion here owes much to Chomsky's treatment of these matters.

67. At least potentially, that is; there is no algorithm for cooking up extensionally equivalent theories.

68. Quine, "Methodological Reflections", p. 106. See also his discussion of Chomsky's notion of tacit rules: "...let me only mention the crying need, at this point, for explicitness of criteria and awareness of method." Ibid., p. 108.

69. See in this regard his "The Problem of Meaning in Linguistics", From a Logical Point of View, pp. 54ff., and his "Reply to Geach" in Words and Objections: Essays on the Work of W.V. Quine, ed. Donald Davidson and Jaako Hintikka, rev. ed. 1975, (The Netherlands: D. Reidel Publishing Company), pp. 328-332.

70. T.G. Bever, J.R. Lackner, and R. Kirk, "The Underlying Structures of Sentences are the Primary Units of Immediate Speech Processing", Perception and Psychophysics, 5 (1969), pp. 225-234. Gilbert Harman has also stressed the relevance of such data in his paper "Quine's Grammar"; this paper helped me clarify a number of issues involved in understanding Quine's views.

71. This, and some of the other questions taken up are not really well-defined at this point, but I don't think that the point made is thereby weakened. Cf. Michael Root, "Quine's Methodological Reflections" in Minnesota Working Papers in Linguistics and Philosophy of Language, ed. L. Hutchinson, M. Kac, and M. Root, September 1973, no. 1. Note that notwithstanding the fact that

there is an indefinitely large 'possible evidence class', we still have no guarantee that our research will lead us to a uniquely correct theory. This 'problem', however, is a feature of empirical research in general, and poses no special difficulties for explanation by ascription of tacit knowledge.

72. Analogously, we might say that the census bureau is interested in population shifts; they are not really interested in why any given individual decided to move.

73. For yet a third non-Chomskyan, non-Quinean view on the status of grammars see Katz, Languages and Other Abstract Objects, (forthcoming).

74. For a discussion of the foundations of this sort of methodological approach to grammars see Katz, Languages.

75. See pp. 11-12.

76. For discussions that seem to pull in both directions at once see his "Philosophical Progress in Language Theory", Metaphilosophy, I, 1, January 1970, pp. 2-19, and his "Linguistics and Philosophy", in Hook, Language and Philosophy, pp. 95-98. The latter includes the following enigmatic comment:

> The as yet unknown innate structures, additional to mere quality space, that are needed in language-learning, are needed specifically to get the child over this great hump that lies beyond ostention or induction. If Chomsky's antiempiricism or antibehaviorism says merely that conditioning is insufficient to explain language-learning, then the doctrine is of a piece with my doctrine of the indeterminacy of translation. (p. 97)

77. "Reply to Harman" (forthcoming).

78. Both Chomsky and Quine seem to agree that ultimate theories will be neither mentalistic nor behavioristic, but physicalistic. The debate is over the placeholder theory we accept in the interim.

79. For an out of character Quinean statement on this issue see the latter part of note 76.

80. The reader will note that I have not discussed these issues in terms of Quine's indeterminacy thesis. Although there are connections, I don't think it would further the arguments to try and make them explicit.

CHAPTER III

CYCLISTS AND SIMULATIONS

The upshot of the last chapter was first, that there is nothing conceptually wrong with the idea that speakers have tacit knowledge of a grammar and second, that the ascription of such tacit knowledge can play a role in explaining a speaker's knowledge of his native language. In this chapter I want to take the defense a step further by showing that we have good reason to think that the speaker's tacit knowledge does in fact play such a role. My strategy remains the same as before; I will be investigating a number of arguments that have been levelled at the Chomskyan conception outlined in Aspects. The difference is that the arguments here all take it for granted that the notion of tacit knowledge makes good sense; the question is only whether it makes good science.

Since these objections and considerations arise from a number of quarters, my presentation may have a slightly scattershot quality to it. In the first section I take up two arguments that straddle the boundary between the a priori issues considered in Chapter II and the empirical matters I want to concentrate on here. Although I ulti-

mately contend that these arguments are not persuasive, they do point to a theoretical gap in Chomsky's original presentation. Much the same point is adumbrated in the next section in the context of a discussion of the status of tacit knowledge vis a vis the distinction between knowing-how and knowing-that. This discussion helps us sharpen our perception of what would be required to fill that gap. In the last section I consider one way developing the notion of tacit knowledge so as to fill this gap in the Chomskyan model.

Two Arguments

Schwartz's objection.

In our discussion of extensional equivalence in Chapter II, I deliberately omitted mention of an argument which might be thought germane to those issues. The intent of that discussion, it will be recalled, was to show that whatever methodological problems concerning extensional equivalence attach to ascriptions of tacit knowledge, apply to ascriptions of explicit knowledge as well. One might have objected, though, that our whole discussion involved an unwarranted assumption. Specifically, it might be asked on what grounds we saw fit to limit the set of extensionally equivalent hypotheses to those involving grammars. Isn't it possible

that there might be one or more hypotheses that account for all the data without attributing tacit knowledge of any rules to the speaker? Consider the following example suggested by Schwartz:

> Suppose we observe an input-output device that labels spheres '+' if their density is greater than 1.0 and '-' if their density is less. We could specify this output with the following equations:
>
> Vol. of sphere=$4/3\pi \text{radius}^3$, density=$\frac{\text{weight}}{\text{volume}}$
>
> Acceptable density (+)=1.0
>
> but the device might never employ a set of principles anything like this. It might never determine the individual sphere's radius, weight or volume; instead it might merely contain a liquid of density 1.0 and label '+' any sphere that sinks in the liquid, and '-' any sphere that floats.[1]

The problem here is not just that the device uses a different set of <u>rules</u> than the one hypothesized. The point is that perhaps a simple mechanical operation having nothing at all to do with rules will explain the input-output correlations well enough. We can even imagine that the labelling itself is carried out 'mechanically', e.g. the displaced liquid weights down a press that prints a '+' or a '-' depending on the amount of pressure exerted.

The objection then is that the same possibility exists in the case of psycholinguistic ascription of tacit knowledge. Even if we grant that the hypothesis of an internally represented grammar could play a role in an explanation of some aspects of linguistic functioning, how do we rule out the possibility that this role is in fact performed without internally represented rules? In terms of the problem of extensional equivalence, the objection will be that we only suggested possible sources of evidence that would discriminate between members of the set of extensionally equivalent grammars; this may not be enough. It could turn out that the real explanation falls completely outside of this set.[2]

My reply to this objection should explain my reason for not including it in the last chapter. In that chapter I was concerned with a priori arguments against the notion of tacit knowledge; this argument doesn't fit into that category. The only conclusion that can be legitimately drawn from the example is that there is no a priori argument to prove that knowledge of language must be explicated in terms of tacit knowledge of grammars. In principle, there will always be other possibilities to consider, and they will all have to be ruled out empirically. I wholeheartedly agree with this conclusion. The ascription of tacit knowledge is an empirical matter through and through. Not only is the 'internal' question

regarding what specific principles, rules, etc. we should ascribe to be decided empirically, but the 'external' question--i.e. whether or not a specific explanandum should be accounted for in terms of underlying tacit knowledge--is also an empirical matter. Still, it is certainly no objection to Chomsky's theory (an empirical hypothesis) that it might be wrong or that it is not true a priori.

Harman's objection

The second argument I want to discuss is presented by Harman in "Psychological Aspects of a Theory of Syntax":

> Taken literally, he (Chomsky) would be saying that we are to explain how it is that Smith knows how to speak and understand a language by citing his knowledge of another more basic language in which he has (unconsciously) "internally represented" the rules of the first language. (It does not seem to make sense to assume that Smith can represent the rules without representing them in some language.) The main problem with such a literal interpretation of these remarks would be the implausibility of the resulting view. How, for example, would Smith understand the more basic language? In order to avoid either an infinite regress or a vicious circle, one would have to suppose that Smith can understand at least one language directly, without unconsciously knowing the rules for that language. But if this is admitted, there is no reason Smith cannot know directly the language he speaks. Thus literally interpreted, Chomsky's theory would almost certainly be false.

Harman goes on to construct a reductio ad absurdum: the assumption that an ordinary speaker internally represents

a grammar leads to the conclusion that some language can be know directly and can serve as the medium for internal representation. But this, Harman argues, "would be proposing that, before he learned any language, Smith had made a presumption about certain data; he had set himself a task; he possessed a theory; he made assumptions; he had techniques of representation. This would be possible only on the absurd assumption that before he learned his first natural language Smith already knew another language."[3, 4]

Two readings. The first thing to notice about this passage is that it really contains two arguments that segue into each other. The reductio that concludes the passage is not so much a spelling out of the conclusion to be drawn from the reasoning that precedes it as it is a new argument:

> (1) Chomsky's use of the notion of tacit knowledge of a grammar implies that speakers have an 'innate' language available.
>
> (2) The notion of such an innate language is absurd.
>
> (3) Therefore, Chomsky's theory of tacit knowledge is absurd.

There is, however, a second, less radical conclusion implicit in the first part of the passage. What the threatened regress of internally represented rules can be taken to show is that there must be some language (or system of representation)[5] that is available to the

speaker even though the speaker does not have tacit knowledge of its grammar. It requires a separate argument to show that this language must be innate. This is what Harman is getting at when he says that at least one language can be known 'directly'. The second argument then is this: if it is not a general truth about knowledge of language that it must be mediated by tacit knowledge of a grammar (i.e. some languages are known directly), then why assume that it holds in the particular case of the speaker's native language? This conclusion is less radical that its predecessor in that it doesn't demand an outright rejection of ascriptions of tacit knowledge. Instead, we are challenged to show why ascription of tacit knowledge is necessary here; why the hypothesis that English, for instance, is known 'directly' will not suffice.

Innate knowledge and direct knowledge. A certain amount of confusion is generated by the failure to distinguish these two very different conclusions. Take, for instance, Chomsky's original response to Harman on this point:

> ...there is a very simple reason why we cannot assume that the child knows the language he speaks directly, without learning, namely, that the assumption is false. We cannot claim that every child is born with a perfect knowledge of English.[6]

Chomsky is no doubt correct in his claim that we are not

all born knowing English, but this is beside the point. This mix-up can be attributed to the fact that he takes 'directly known language' to be another way of talking about 'innately known language'. But for Harman this is an empirical hypothesis that requires further evidence. To know a language 'directly (for Harman) is not to be born with it or to know it 'without learning' as Chomsky imagines. It is simply to know it without first coming to have tacit knowledge of its grammar. Chomsky's reply, then, assumes that if a speaker knows a language, either (a) he knows it directly--i.e. the language is innate in the speaker, or (b) he knows it indirectly--i.e. it was acquired by internalizing the rules of its grammar. As a response to the first conclusion we distinguished, the reply is wide of the mark: Harman finds the whole idea of an innate language 'absurd'. He certainly does not hold that English is innate. As a response to the second argument, the reply begs the very question Harman wants to pose: if some language can be known directly, how do we rule out the possibility that a speaker's native language is so acquired?

We could go on to carefully dissect the assorted misunderstandings that surround this argument, but such exegetical concerns are not to the point. It is worth noting, however, that as far as the first argument is concerned, Harman concedes that what he originally took

to be a reductio is in fact a sound inference:

> I was wrong when I said that Chomsky's theory about how someone (Smith) learns a language "would be possible only on the absurd assumption that before he learned his first natural language Smith already knew another language." Not only is the assumption not "absurd", it is probably true in the relevant sense of "language". For Smith can probably think before he learns his first natural language, and thought requires a system of representation, where any such system must count as a language in the relevant sense of "language".[7]

But since he goes on to deny that "such a concession affects the force of my objection to Chomsky's account of language learning," it is safe to assume that the key question is the one about direct knowledge of a language: viz. if we are not compelled to explain knowledge of language in terms of tacit knowledge of grammars, then why should we?

Reply to Harman. As it stands, I think this last argument can be countered in much the same way we dealt with the earlier objection suggested by Schwartz's example. The general import of both arguments is as follows. There are capacities (psychological and other) that can be described in terms of rules/principles that fit their inputs and outputs. Nevertheless, it does not follow that the exercise of these capacities involves the application of these or any other rules. Schwartz's case provides a convincing example of this principle in ac-

tion; Harman establishes the same point on the basis of a philosophical argument showing that there must be some rule-fitting operations that are not cases of rule-guidance. Neither argument presents a bona fide objection to Chomsky's conception of tacit knowledge. A bona fide empirical objection would involve showing that in fact there is an explanatory framework that accounts for linguistic functioning more simply than the Chomskyan framework which makes the assumption that speakers have such tacit knowledge of a grammar. A proof that such alternative frameworks are possible is not a sound reason for abandoning, or for that matter even losing confidence in, a model that has hopes of success.

Rules and explanations

Despite the fact that I take the foregoing reply to be correct, I think that there is a critical insight to be captured that these arguments barely miss. Perhaps we can begin to see what this insight is by considering the following moves that might be made at this point.

Someone might say that this talk of 'alternative explanatory frameworks' introduces an unnecessary complexity. Put as starkly as possible: the explanandum here is the set of phenomena we colloquially refer to as 'knowing a language'. We are faced with a choice; we can say (1) that the speaker acts in accordance with the

rules of the grammar, or (2) that the speaker follows the rules which he tacitly knows. Perhaps the real objection to be made is that (1) is explanation enough. In other words, the simple empirically supported claim that such-and-such aspects of linguistic functioning obey certain regularities is already an explanation of 'knowing a language'. For those who identify explanatory adequacy with predictive power, this reply will be especially attractive. Defenders of (1) have simplicity on their side in that they make the minimal claim; what additional grounds are there for supporting (2)?

What this imaginary objector says to Chomsky is this: "Your theoretical work in linguistics is wonderful. The results so far are insightful and capture a large number of significant and interesting regularities. Still, there is no reason to take a perfectly satisfactory theory of how the mind works and say that the mind applies[8] this theory in its own workings. Such a step is unnecessarily extravagant."[9]

Suppose we reply that only (2) provides an explanation. (1), we say, has no explanatory value at all; it only gives us a perspicuously clear summary of the explanandum. In saying that the speaker acts in accordance with such-and-such rules, we have simply organized the phenomena that have to be explained--we have not yet ex-

plained anything. We can still ask why it is that the speaker acts in accordance with this set of rules. What is special about the speaker that explains the peculiarities and regularities of his knowledge of language? It is not enough to reply that his knowledge of language involves these peculiarities and regularities.

This is not the final word, however. Our opponent will argue that our standards of explanatory adequacy are too high. To explain the behavior of the falling rock in Chapter I, isn't it enough to point out that rocks fall in accordance with well-known laws and that these laws entail that this rock will fall in just the way it is observed to fall? Must we add that the rock is following some rule/law that is internally represented within it? Certainly not. Why then should we feel a need to say the speaker does? What is it about knowledge of language that makes this extra step appropriate?[10]

The foregoing conceptual jousting has been lurking in the background of all our discussions thus far. I think we are beginning to get closer to the insight that lies behind many of the empirical challenges brought up earlier. At this point I want to introduce one more issue to help sharpen the question. This issue concerns the relation between tacit knowledge, on the one hand, and the distinction between knowing-how and knowing-that

on the other.

Knowing-how and Knowing-that

Harman's diagnosis

Apart from the specific criticisms of Chomsky's position that Harman makes in the previously cited papers, he also attempts to diagnose the basic 'philosophical' error of the Aspects program. As Harman sees it, the wrong-headedness of the whole methodological approach can be traced to Chomsky's failure to distinguish knowing-how from knowing-that:

> I believe that Chomsky's notion of linguistic competence is best seen as the result of two separate but interacting factors. First, chomsky's use of the phrase 'tacit competence' betrays a confusion between the two sorts of knowledge of a language. Competence is knowledge in the sense of knowing how to do something; it is ability. It is not the sort of thing that can properly be described as "tacit". Tacit knowledge must be knowledge that something is the case. Second, and this abets the above confusion, speakers of a language do have something that might be thought of as tacit knowledge about the language. Thus, speakers can be brought to judge that certain sentences are ambiguous, that certain sentences are paraphrases of each other, or that certain strings of words are not grammatically acceptable.[11]

I want to put aside Harman's second factor for the time being since it will be discussed more fully in the next chapter. Here I want to get clear about the first.[12]

This criticism of Chomsky--i.e. that he confuses knowing-how and knowing-that--operates on two different levels that happen to run parallel to two different ways of understanding the phrase 'knowledge of language'. When we refer to someone as knowing English, for example, we usually mean that the person has (at least) a whole network of abilities--for instance: the ability to express himself in English, to understand other typical English speakers, to spot flagrant misuses of the language, and so on. What Chomsky is trying to do, as a psychologist-linguist interested in a particular language, is to provide a theory that will explain this knowledge. One might say that he is trying to fill in the schema 'Knowledge of English is a matter of X'. The competence/performance model, and the attendant notion of tacit knowledge of a grammar provides a partial theory of how we are to replace the 'X' in the above schema. If Chomsky is right, the knowledge of English is really a matter of having tacit knowledge of a grammar. That is to say, the crucial (but not sole) underlying cause of what we colloquially refer to as 'knowledge of English' is this tacit knowledge.[13] Ultimately, it is in virtue of their knowledge of language in this second sense that speakers have knowledge of language in the first colloquial sense.

When Harman charges that knowledge of language is a matter of knowing-how, not knowing-that, he could be taken to mean one or both of the following: (a) that when we say 'A knows language L' we do not mean to assert that A has a body of propositional knowledge at his fingertips, but that he has a certain set of abilities; (b) that the underlying mental/brain structures that are responsible for a speaker's knowledge of English are not to be characterized in terms of propositional knowledge of rules, principles, etc.--i.e. that there is no real tacit knowledge that underlies these abilities.

If Harman means (a), then there is no real problem. He is probably right in his contention that when we say someone knows English we don't mean that he has a body of truths at hand that non-English speakers lack. Although he may very well have such a body of truths available, that is still not what we mean when we say he knows English. Be that as it may, I can find no evidence that Chomsky has ever challenged this view, and it is not incompatible with the explicit program laid out in Aspects. The second criticism (b) is more interesting and seems to mark off a real area of disagreement. Here Harman seems to defend the irreducibility of know-how to know-that; that knowing how to do something is not a matter of having a body of underlying propositional knowledge (such as knowledge of a grammar). It is not that Harman has

qualms about taking a grammar as a fit subject of propositional knowledge. I think that what lies behind this claim is a Wittgensteinean point that is stated forcefully by Ryle.[14]

Ryle's objection. In his discussion of the relation between the behavior of a reasoner and the deductive principles formulated by logicians, Ryle says this:

> There is a not unfashionable shuffle that tries to circumvent these considerations by saying that the intelligent reasoner who has not been taught logic knows the logicians' formulae 'implicitly' but not 'explicitly'... This shuffle assumes that knowledge-how must be reducible to knowledge-that, while conceding that no operations of acknowledging-that need be actually found occurring. It fails to explain how, even if such acknowledgments did occur, their maker might still be a fool in his performance.
>
> All this intellectualist legend must be rejected, not merely because it tells psychological myths but because the myths are not of the right type to account for the facts which they are invented to explain. However many strata of knowledge-that are postulated, the same crux always recurs that a fool might have all that knowledge without knowing how to perform, and a sensible or cunning person might know how to perform who had not been introduced to those postulated facts; that is, there still remains the same gulf, as wide as ever, between having the postulated knowledge of those facts and knowing how to use or apply it; between acknowledging principles in thought and applying them in action.[15]

As applied to our case, the point would be that ascribing a set of internally represented rules is theoretically useless since such an ascription would not help us ex-

plain the speaker's know-how.

Reply to Ryle. I think that Ryle is right in that even if we do postulate some tacit knowledge of a grammar, we will still have to explain how it is that this knowledge is effective in bringing about behavior. This problem goes beyond the issue of what role a competence component plays in a performance model of linguistic behavior; it addresses a far deeper philosophical question: how can knowledge, which we think of as being in some sense static, affect behavior, which is prototypically dynamic?[16]

I think that Ryle is wrong, however, in attacking the supposition of implicit knowledge because it offers no solution to this problem.[17] At least he is wrong in thinking that even with such a supposition the explanatory gulf is "as wide as ever". Consider the following example (which recalls Quine's Danish schoolboys). Imagine that a student is taught the syntactic rules of the propositional calculus (PC). Assuming that he has learned his lessons, he can represent English sentences in PC, he can pick out well-formed formulas of PC, and so on. It is natural to account for many of his PC-related abilities by citing the fact that he was taught the syntactic rules and that he knows how to apply them. Of course the account would be incomplete if we left out the

last clause. If we simply said that he was taught the rules (or even that he knows them, that he has formed an internal representation of them, etc.) that would not be enough. Now the Rylean is certainly right in fixing on the phrase 'and he knows how to apply them' as requiring further unpacking. But so what? How it is that knowledge--tacit or explicit--gets applied is a general problem in the philosophy of mind. We have been given no reason to think that, if and when this problem is solved, the solution will not also explain how our tacit knowledge is applied. Is the suggestion here that we balk at all explanations of behavior in terms of knowledge--e.g. that Cohen ran for the train because he knew that he was late--on the grounds that it leaves unexplained the relation of knowing and behavior?

As I see it, in trading the question "How does the student manage to pick out the well-formed formulas of PC?" for "How does the student (know how to) apply the rules he knows?" we are trading up; we are making progress in our attempt to understand his behavior. We must be ready to concede that the question we are left with is extremely difficult. We might ultimately conclude that some knowledge just 'gets applied' without the subject's knowing how to apply it. Perhaps at a deep enough level we will be satisfied with an explanation of certain types of applications of knowledge in terms of physiological

mechanisms. It might even turn out (horror of horrors) that we will never understand how the knowledge of the syntactic rules of PC enables the student to perform the discriminatory task. Nevertheless, the fact that knowledge-based explanations have to stop somewhere does not imply that they ought not to start anywhere.

My response to the Rylean objection is certainly not new, but I think it is on target. The current interest in cognitive psychology attests to the fact that this particular worry of Ryle's has not carried the day. Perhaps the positive lesson is that ascription of tacit knowledge of grammar cannot be the <u>whole</u> explanation of knowledge of language; there is more to be said about how the explanatory model might be supplemented. But this is more of a reminder than a lesson--no one, to my knowledge, has claimed that the application of the speaker's tacit knowledge of a grammar is self-explanatory. As far as Harman is concerned, we will say that even if we do want to characterize knowledge of a language in terms of knowing-how, it has not been shown that it cannot be explained--at least in part--by positing an underlying base of tacit knowledge of rules.

An aside: we might note in passing that a pattern of attack is emerging. A significant number of the objections to the notion of tacit knowledge that we will

consider turn out to be problems about knowledge and belief in general. Ryle's objection here and the problem of extensionally equivalent grammars of the previous chapter are typical cases of this sort of misdirected attack. There is a tendency, then, to take the modifier 'tacit' to signal some special sort of vulnerability. I have assiduously argued that this is a mistake.

The problem of the cyclist

The reply to Ryle, however, suggests a possible way to press the point against Chomsky. If knowledge of language is indeed a case of knowing-how and is to be explained in terms of underlying tacit knowledge of a grammar, then why not apply the same explanatory paradigm to all cases of knowing-how? Harman makes this point in terms of knowing how to ride a bicycle.

> The cyclist must keep balanced on the bicycle. Exactly what he needs to do in order to retain his balance will be dictated by certain principles of mechanics that he himself is unaware of. We might imagine that a model of a cyclist would contain representations of the relevant principles of mechanics and that it uses these principles in calculating what needs to be done so as to retain balance. Then, following Chomsky's model, we might say that the cyclist has an internal representation of the principles of mechanics. Should we go on to say that every cyclist has an intuitive or tacit knowledge of the principles of mechanics?[18]

Harman takes it as obvious that the answer to his last question is 'no'. He goes on to say that attributing

knowledge in this case "does not seem to be an illuminating way of talking about cyclists; and Chomsky's remarks about "tacit competence" do not seem to provide an illuminating way of talking about speakers of a language."

I take this question about the cyclist to be a key challenge to Chomsky's use of the notion of tacit knowledge; this is exactly where there is a gap in the original presentation in Aspects. Unless we can constrain ascriptions of tacit knowledge--that is, unless we can construct criteria that rule out ascription of tacit knowledge in the cyclist's case but not in the speaker's case--we will have an explanatory tool that is too powerful. To put it slightly differently, we will end up taking as rule-guided behavior what is really only rule-conforming behavior.

Chomsky's response

In dealing with this objection there is a great temptation to backtrack. Perhaps we were too quick to grant Harman's point that knowing a language is a matter of knowing-how. In his reply to Harman, Chomsky has this to say about the cyclist argument:

> ...he (Harman) suggests no respect in which ability to use a language (let alone the competence, in my sense, that constitutes an element of this ability) is like the ability to ride a bicycle, nor do I see any. The proper conclusion, then, would be that there is no

> reason to suppose that knowledge of language can be characterized in terms of "knowing how". I therefore see no point in the analogy that he suggests. Knowledge of language is not a skill, a set of habits, or anything of the sort. I see nothing surprising in the conclusion that knowledge of language cannot be discussed in any useful or informative way in this impoverished framework. In general it does not seem to me true that the concepts "knowing how" and "knowing that" constitute exhaustive categories for the analysis of knowledge.[19, 20]

The intention here is to undercut the objection by attacking one of its premisses. If knowledge of language is not know-how, then the comparison of the speaker and the cyclist will seem beside the point. Since Chomsky also claims earlier that knowledge of language is not a matter of knowing-that, he is driven to reject the how/that dichotomy as 'impoverished'.

I think Chomsky may be right about this last point. Perhaps we need another category of knowledge--neither knowing-how nor knowing-that--and perhaps knowledge of a language falls under this third category.[21] Chomsky's conception of psychological research as the discovery and appropriate characterization of human competences or cognitive capacities seems very much in line with this tripartite conception of knowledge.[22] Be that as it may, pointing out the inadequacies of the division of knowledge into knowing-how and knowing-that does not provide us with a solution to the cyclist problem. The problem

is only suggested by the view that knowledge of language is a matter of knowing-how. It will not go away even if we give up this assumption. It is still incumbent on those who want to defend ascriptions of tacit knowledge to explain in what way the speaker's case differs from the cyclist's. Chomsky's response here boils down to the claim that the two are different because knowledge of language involves competence--i.e. tacit knowledge of a grammar--while knowing how to ride a bicycle does not. But this puts the cart before the horse. The challenge is to show why we should not also explicate knowing how to ride a bicycle in terms of a specialized intellectual competence--i.e. tacit knowledge of the relevant principles of mechanics. Even if we want to classify knowledge of language in a separate category from instances of ordinary know-how, we will still have to present criteria for making these distinctions: how do we determine which abilities are to be explicated in terms of underlying tacit knowledge and which are not? So even if we could mount a successful attack against the know-how/know-that distinction, we would still not have met the challenge of the cyclist head on.

In subsequent discussion of this issue, Chomsky comes to see that his original attack on the how/that distinction is, when considered in isolation, an insufficient response. The fundamental point, then, must be

that the ascription to the speaker can be justified as an instance of inference to the best explanation, while the parallel ascription to the cyclist cannot:

> ...we do not attribute knowledge of mechanics to he bicycle rider if in fact this assumption does not help explain his ability to ride a bicycle; we do attribute knowledge of the rules of grammar to the speaker-hearer if this assumption does contribute to an explanation of his ability to use a language.[23]

But this response is still too vague. Why would the ascription to the cyclist not play a role in explaining his ability to ride? Couldn't we develop a story to the effect that the cyclist wants to turn a corner, has a certain perceptual input p, converts the input into a direction for his muscles by applying the principles of mechanics he knows tacitly, etc., and eventually turns the corner?

It might be thought that Chomsky has the following distinction in mind: ascription of tacit knowledge of the principles of mechanics would not help explain the cyclist's ability because there are plenty of physicists who have explicit knowledge of the principles of mechanics and are still unable to ride a bicycle. This, however, is not a response available to Chomsky. This is simply the analogue of the Rylean argument we rejected earlier; it does not provide any distinction between the cyclist and the speaker. There may very well be archae-

ologists who have explicit knowledge of the grammar of Linear B who don't know the language in any sense relevant to our discussion. As I argued, this does not indicate that knowing a language does not involve tacit knowledge of rules, but only that tacit knowledge of rules is not sufficient for knowing a language. The ascription of tacit knowledge to both the speaker and the cyclist would in both cases leave us with only incomplete explanations of the respective abilities. Harman's challenge still stands then.

Chomsky's notion of tacit knowledge can be developed in two different directions at this point, depending on how we treat Harman's cyclist argument. One approach would be to deny Harman's claim that the ascription of tacit knowledge of mechanics "does not seem to be an illuminating way of talking about cyclists". Following this line we would argue that notwithstanding the initial implausibility attached to this claim about cyclists, the best way to explain the cyclist's ability to keep his balance involves ascription of such underlying knowledge to him. This is very like the approach taken by Fodor in his treatment of tacit knowledge. A second approach will be to accept the challenge of Harman's cyclist argument but to argue that there is a significant difference between the speaker and the cyclist that warrants the ascription of tacit knowledge to the former, but not to

the latter. This strategy is developed by Graves, Katz, et. al.. In what follows I want to look closely at these two ways to develop the notion of tacit knowledge. Here I take up Fodor's theory; we will turn to Graves, Katz, et. al. in Chapter IV.

Fodor on Tacit Knowledge

Let us begin by reviewing the Rylean argument taken up earlier. Ryle, we recall, would discredit Chomsky's tacit knowledge strategy by associating it with 'the intellectualist legend'. The reasoning, again, is this. (1) All intellectualist explanations of know how in terms of underlying propositional knowledge are inadequate; (2) ascriptions of tacit knowledge of a grammar are instances of such explanations; so (3) ascriptions of tacit knowledge of a grammar are inadequate. Fodor's theory can be understood as a direct counterattack against Ryle's argument. He sets out to defend intellectualist explanations in general, and tries to accord the concept of tacit knowledge a systematic place in all such accounts. For Fodor, tacit knowledge is, as it were, woven into the fabric of psychological theorizing. His aim, as he puts it, is to "defend 'intellectualist' accounts of mental competences and to suggest that, in attributing tacit knowledge to organisms, intellectualist theories exploit a legitimate form of nondemonstrative inference:

the inference from like effects to like causes".[24] His argument begins with a theory of psychological explanation and ends with a suggested set of criteria for ascribing tacit knowledge on the basis of such a psychological theory. In what follows I want to first sketch out the basic structure of Fodor's position and then see whether the connecting links he provides give us a sound argument. We begin with his theory of psychological explanation.

Fodor's argument

Fodor on psychological explanation. For Fodor, a psychological theory is a set of explanations of the behavioral capacities of the organism. Each explanation takes the form of an answer to the question: How does the organism produce behaviors of type t? The prototypal answer to such a question will be the specification of a sequence of instructions for carrying out the behavior. In the 'completed theory', each of the instructions in the sequence will be 'elementary'--that is, each will be an instruction to carry out what Fodor calls an 'elementary operation'. He considers an operation elementary if it is "one which the normal nervous system can perform but of which it cannot perform a proper part".[25] What we have then is a compositional theory of human behavior in

which complex behaviors are functions of simpler elements in the organism's repertoire. The theory achieves its explanatory power not by predicting the occurrence of any given behavior, but by accounting for the multiplicity of behavioral capacities in terms of the relatively small set of elementary operations and the set of possible combinations of such operations. Each explanation can be thought of as a recipe for producing behaviors of a certain sort.

This idea of combinations of elementary instructions suggests a computer program as an appropriate simulation model for psychological processes. Fodor casts the analogy in these terms:

> .. every box name in a computer flow chart is the name of a problem. If the computer is to simulate behavior, every box name will be the name of a psychological problem. The problem is to specify a sequence of instructions which
>
> > (a) will convert the input to the box into the output from the box, and
> >
> > (b) can be written in a way that mentions only operations that are elementary for the organism whose behavior we are trying to simulate.
>
> The problem is solved (an optimal simulation is achieved) when instructions are formulated which satisfy both these conditions and which meet the usual methodological constraints upon theories: simplicity, conservatism, coherence with theories in related sciences, etc..[26]

To say that a theory based on such an optimal simulation is true--i.e that the optimal simulation is correct--is

to hold that the organism's behavior is mediated by the same sequence of elementary operations that is specified in the program.[27] I will not try to say here what I think the ultimate cash-value of the computer analogy is. At the very least, devising a simulation in these terms provides a convenient way to state the details of a theory and to test its specific consequences.

The general picture of psychological explanation that emerges from this approach is captured in Fodor's celebrated 'explanation' of how we tie our shoes:

> There is a little man who lives in one's head. The little man keeps a library. When one acts upon the intention to tie one's shoes, the little man fetches down a volume entitled Tying One's Shoes. The volume says such things as: "Take the left free end of the shoelace in the left hand. Cross the left free end of the shoelace over the right free end of the shoelace... etc."
>
> When the little man reads the instruction 'take the left free end of the shoelace in the left hand', he pushes a button on a control panel. The button is marked 'take the left free end of a shoelace in the left hand'. When depressed it activates a series of wheels, cogs, levers, and hydraulic mechanisms. As a causal consequence of the functioning of these mechanisms, one's left hand comes to seize the appropriate end of the shoelace. Similarly, mutatis mutandis, for the rest of the instructions.
>
> The instructions end with the word 'end'. When the little man reads the word 'end', he returns the book of instructions to his library.
>
> This is the way we tie our shoes.[28]

Whatever its educational value, it should be clear that this picture of psychological processing in terms of a supervising homunculus who consults manuals of instructions is the Rylean nightmare come alive. It was the fear that 'this is just what intellectualist explanations will come to' that motivated the Rylean attack.

But where does this story get us? The fact that we can tell an anthropomorphic fairy tale about an organism's internal processing and talk as if armies of little men were following instructions does not really pose a serious problem for anti-intellectualists. No one wants to deny that there are micro-processes involved in shoe tying that are beyond the pale of our awareness. The anti-intellectualist will also grant, at least for the sake of argument, that the central nervous system can be analyzed in terms of a set of elementary operations, and that complex behavior is the product of coordinated structures of such operations.[29] As long as we keep in mind that we <u>are</u> anthropomorphizing--that there are no little men--why should a Rylean balk?

Fodor's position here, as best I can make it out, is that if we look to current experimental results in psychology, we are forced to accept intellectualist explanations. The example he provides has to do with perceptual judgements of depth. In this regard, he points out

that the according to the best psychological evidence available, human depth estimation is mediated by the perception of texture gradients. Starting with this empirical fact, Fodor argues that

> . . . if the central nervous system is the kind of organ most sensible people now think it is, then some of the things the central nervous system _does_, some of the physical transactions that take place in the central nervous system when we make estimates of depth, must satisfy such descriptions as 'monitoring texture gradients', 'computing derivatives of texture gradients', etc. ... whatever physical system mediates our perception of depth must, _ipso facto_, perform whatever operations are nomologically necessary and sufficient for depth perception. If, then, our perceptions of depth vary as a function of the first derivative of texture gradients, then among the operations of the central nervous system must be some which satisfy such descriptions as 'is computing the first derivative of texture gradient _t_'.[30]

The evidence is supposed to show that the causal connection between the environment and our judgements of depth are mediated by operations that can be truly characterized in terms of monitoring, computing, and so on. If we accept these results (and we have no experimental reason not to), then we are _ipso facto_ committed to an intellectualist explanation. Computing the first derivative of a texture gradient is an intellectual operation that involves the application of computational principles. Such principles, then, comprise part of the intellectual substratum of depth perception (and estimation). The final step in the argument involves the claim

that if a person's behavior is in fact mediated by rules and principles of this sort, and he has no awareness of applying such principles, then we should say that he knows them tacitly.

Fodor's criteria. On the basis of the foregoing, Fodor recommends that we accept the following criteria for tacit knowledge ascription.

> . . . if X is something an organism knows how to do but is unable to explain how to do, and if S is some sequence of operations, the specification of which would constitute an answer to the question "How do you X?" and if the optimal simulation of the behavior of the organism X-s by running through the sequence of operations specified by S, then the organism tacitly knows the answer to the question "How do you X?", and S is a formulation of the organism's tacit knowledge.[31]

We can (perhaps artificially) divide this set of criteria into two parts--general conditions for ascription of tacit knowledge, and criteria for characterizing the specific content of the tacit knowledge. The conditions that determine an appropriate context for ascription are simple: (1) the organism knows how to do some X, but (2) cannot explain how it does X. Once we've determined that a given bit of behavior should be explained in terms of the tacit knowledge of the organism (that we are dealing with what might be called an 'intentional context'), we turn to our psychological theory: (3) we ascribe tacit

knowledge of whatever set of instructions an optimal simulation of the organism would execute to produce type-X behaviors.

Problems for Fodor

I eventually want to argue that this conception of tacit knowledge is confused; that it is the result of running together two incompatible views as to the nature and justification of intellectualist explanations. We can begin to see how this confusion arises by taking a closer look at Fodor's criteria. To begin with, we might wonder what Fodor has in mind in delimiting ascriptions of tacit knowledge to cases where the organism can't 'explain how' he does what he 'knows how' to do. Both these concepts need unpacking.

Knowing how. The problems in spelling out what is to count as know how are notorious. Even though Fodor cannot be blamed for this state of affairs, we nevertheless have to know what he intends to include under this rubric if we are to be able to apply his criteria. Even on a conservative reading, all of the following can be counted as instances of know-how:

knowing how to soothe a baby's crying

knowing how to get your crying soothed

knowing how to prove a theorem

knowing how to wiggle your ears

knowing how to ride a bicycle

knowing how to speak and understand English,

Does Fodor think that all these cases are candidates for tacit knowledge explanation? One is tempted to say that unless we have reason to think otherwise we ought to assume so--i.e that all the above satisfy the first condition. But on the other hand, it is not obvious that these criteria fit Fodor's earlier stated aim--i.e. "to defend intellectualist accounts of mental competences" (Italics mine). Does the intellectualist want to say that knowing how to wiggle one's ears is a mental competence? That knowing how to get your mother to soothe your crying (cry louder?) is a mental competence that babies have?

Explaining how. Turning to Fodor's talk of 'explaining how' we encounter further difficulties. The intent in including this condition is innocent enough: Fodor does not want his criteria to lead to the ascription of tacit knowledge of some set of instructions S to an organism that already knows S explicitly. Nevertheless, there are are two problems that should be noted.

First off, 'explaining how' is not much clearer than 'knowing how'. There are any number of cases where we would be hard put to say whether what the subject offers

as an account of his behavioral capacities should count as 'being able to explain how he does it'. Do we 'explain' how we perceive depth when we say "I just open my eyes and look and I know right off."? In one sense, this is how we do it. In another sense, however, we can be genuinely awed by our ability to judge depth as something we haven't the faintest explanation for. In this second sense, "I just open my eyes . . ." is no more than a reconstruction of the preliminaries; it is not an explanation of the real thing. Although it is not clear how we are to draw this distinction, it might be assumed that there is a distinction here and that Fodor must have this second sense in mind. For one thing, if he did not, the depth perception example would be completely inappropriate.

I want to argue, however, that the conception of tacit knowledge ascription that emerges from this second sense of 'explaining how' is incongruous with the criteria that Fodor finally suggests. To see why this is so, we have to consider what Fodor's ascriptions of tacit knowledge are going to look like. Fodor is committed to the view that what we ascribe is knowledge of some S that an optimal simulation runs through in producing the relevant behavior. But by the definition of optimal simulation, S is a sequence of elementary instructions. So it would seem that what the organism tacitly knows,

putting it roughly, is that to bring about some desired behavior b it can/must execute a sequence of elementary instructions S. This is what it means to say that the sequence is the 'recipe' for bringing about the desired behavior. As examples of the types of operations/instructions Fodor has in mind we have monitoring texture gradients, computing first derivatives of texture gradients, and so on. (These are probably not going to be the ultimate elementary operations that the nervous system is wired to perform, but we can assume for the sake of argument that they are.)

The point I want to make is that given this view of an optimal simulation and the constraints it puts on what a psychological theory can accept as an answer to the question 'How do you X?', it is extremely far-fetched to think that anyone can really explain how they do anything. Our everyday explanations of know-how are in no sense basic--they are not couched in terms of the elementary operations in our behavioral repertoire. What we try to do when we explain how to throw a curveball, for instance, is to break down a complex behavioral competence into a series of more elementary behavioral competences--"You grip the ball like this and then break your wrist as you whip your arm around". This looks very much like what Fodor wants a psychological explanation to do. The crucial difference however, is that Fodor's psycho-

logical explanations do not stop at the level of elementary behavioral competences. There is every reason to think that although we can't get more nitty-gritty than 'whip your arm around like so', a psychological explanation in terms of elementary operations will go deeper. The computational operations and instructions that will appear in this deeper explanation will no doubt lead to ascriptions of tacit knowledge that even the most discursive pitching coach will be unaware of. So I think that Fodor is committed to the view that in the sense of 'explaining how' that is relevant to the intellectualist--the sense in which a psychological theory explains how people are able to do what they can do-- no one can ever really explain how they manifest their know-how.[32]

To forestall any misunderstanding, let me say that I don't mean to recommend a new condition on explanation to the effect that every explanation of know-how that is not framed in terms of elementary instructions is bogus. Explanations must be suited to their context, and it is difficult to imagine many contexts where we would want explanations in terms of elementary operations. This is all true but beside the point. The point is that no matter how much explaining the 'organism'[33] does--no matter how much of his explicit knowledge he can marshal--the consistent intellectualist (who adopts Fodor's model) will have to say that at bottom the behavior is in

fact mediated by the organism's tacit knowledge, and that this knowledge has the form of a set of elementary instructions for producing the behavior.

In itself, this is not an especially devestating criticism of Fodor's theory. He could simply admit that the explaining-how condition ((2)) is dispensible, and claim that all our behavior, whether we can explain its etiology or not, is mediated by our underlying tacit knowledge. What I think is important about this criticism is that it points to the fact that there are two different conceptions of tacit knowledge operating in Fodor's model. I want to try to trace through these different conceptions and explore their respective consequences.

The argument from simulation and the argument from computation. We can start by noticing that Fodor's premisses can be called upon to support either of two completely different arguments for ascribing tacit knowledge. Both proceed from like effects to like causes, and both make use of the notion of an optimal simulation. The arguments can be stated in this way:

The argument from simulation (AS). An optimal simulation of an organism produces the same set of 'behaviors' as does the organism, and is equipped with the same repertoire of elementary operations. The operations that a simulation runs through in producing its output are mediated by internally represented instructions (in the form of a program). These instructions, then, constitute an answer to the question 'How do you X?'. It is therefore a reasonable inductive inference from like effects to like causes to say that the organism has also internally represented a set of instructions which it applies in producing behaviors of type X. So the organism also has the answer to the question 'How do you X?' but cannot get at that answer by introspection. We should say, then, that the organism's knowledge is tacit.

The argument from computation (AC). An optimal simulation of an organism produces the same set of 'behaviors' as the organism does and performs only operations that the organism can perform. If cognitive psychologists are right, then some of these operations involve such processes as computing first derivatives. This gives us a strong inductive basis for saying that the organism being simulated also carries out such computational processes. But we ordinarily explain the ability to perform this sort of complex mental operation by positing some underlying knowledge of general prin-

ciples. We should therefore use the same explanation here; except that in this case, since the knowledge is not introspectible, it must be tacit.

Later I will contend that both of these arguments fail to secure their conclusion: i.e. that organisms have tacit knowledge. Before going on to that, however, I want to note two crucial differences between AS and AC.

AS and AC compared. (1) First, notice that AC's plausibility is dependent on the specific empirical theories that cognitive psychologists have put forward, while AS's does not. I don't mean that AC rises and falls with every modification in detail; still, a significant change in psychology would rob it of much of its force. Let us imagine that the cognitivist theory of depth estimation proves in the end to be unfeasible. Imagine further, for the sake of argument, that we discover a simpler explanation along lines closer to standard behaviorist theory. If all we could infer from this hypothetical 'behaviorist' simulation was that certain simple associative operations mediated the stimulus-response nexus,[34] then there might be no grounds for AC's conclusion that our behavior is mediated by unconscious mental operations. AS, on the other hand, derives its result (invalidly, I will argue) on the basis of purely methodological considerations about the general structure of psychological explanations. Since in the simulation

the output involves internally represented recipes, we too must have our internally represented recipes. It is of no matter what sorts of instructions are contained in these sequences, because the argument is not based on the nature of the elementary operations carried out, but on the mere fact that they are carried out. If they are carried out, then we must know some program of instructions for carrying them out.

(2) This leads us to the second important difference. According to AS what the organism tacitly knows in every case is a recipe for producing behavior; that is, the organism knows some mobilization plan for organizing elementary operations so as to produce some complex bit of know how. This is the sort of tacit knowledge that can count as an answer to the question 'How do you X?'. For AC, however, what the organism tacitly knows is not some recipe for producing behavior, but rather a set of rules or principles for carrying out a procedure for reaching a result; in a sense, the tacit knowledge provides a transition between representational states. In the case of depth perception, for instance, this 'result' has the form of an estimate of distance. These characterizations of the difference between AS and AC in terms of 'results' 'behavior', and 'representational states' leaves much to be desired, but we will say more about this dichotomy in the next chapter. For now

it is important to see that AS takes the organism's tacit knowledge to be instructions for bringing about behavior; AC takes tacit knowledge to consist in (computational?) principles that mediate derivations.

It goes without saying that the tacit knowledge AC would have us ascribe will, like all knowledge, have its role to play in the production of behavior. The difference is only one of degree. If you didn't know (tacitly) principles for computing the first derivative of a texture gradient, then you can't compute first derivatives. If you can't do <u>that</u>, then you can't estimate how far you are from objects in your visual field. You are now bound to bump into things quite often. So we can explain Linda's knowing how to get around her living room without doing too much damage by citing her tacit knowledge of computational principles which enable her to compute first derivatives. But here the tacit knowledge plays an indirect role: it mediates the transition between the first two steps of this chain. Contrast this with the tacit knowledge AS would ascribe, which would be more like a game plan for supervising the whole sequence--viz. knowing how to get around your living room.[35]

As we can see, the tension in Fodor's presentation between AS and AC makes it difficult to determine what Fodor wants to ascribe as tacitly known in any specific instance--e.g. the cyclist case. Generally speaking AS is closer to the explicit criteria Fodor recommends which incorporate the idea that tacit knowledge is of some sequence of instructions that would count as an answer to the question 'How do you X?'. On this view, what we know tacitly is a set of 'directions' for bringing about behavior. If we look not at Fodor's stated criteria but instead to the specific considerations that are supposed to lend plausibility to these criteria, it seems that AC is the argument he has in mind. This would explain the introduction of the evidence about the underlying computations involved in depth estimation and the stress on the goal of explaining mental competences.

Problems for AS and AC. Whatever Fodor actually has/had in mind, we are left with two arguments that might be taken to show that organisms have tacit knowledge. I want to show that neither one works.

The case for AS. This argument, as it stands, is especially unconvincing. The mere fact that an optimal simulation performs because of the instructions represented within it, does not show that the simulatee also

performs on the basis of internally represented instructions. Consider in this regard a meteorologist's simulation of Hurricane Hannah. The device will only run through the processes it does because of its internalized program of instructions. Still, even the staunchest meteorological realist will not conclude that the hurricane has any sort of internally represented instructions.

We construct a simulation to develop a hypothesis about the underlying features of the simulatee's behavior. The program of instructions is simply the way the machine performs (or: outputs descriptions of) these operations. Unless we have a separate argument to the contrary, the fact that the simulation uses a program is no more relevant to our theory of the simulatee than is the fact that the program is stored on a special sort of disc. All the optimal simulation can do is give us evidence concerning the specific operations that are performed in mediating the organism's behavior--e.g. that texture gradients have to be monitored. It can play this role because we have access to the program which is a record of its mediating operations. So AS's argument from the fact that operations are performed in the organism to the claim that there is some internalized program that the organism tacitly knows, does not go through.

Considering my prior attacks on Schwartz and Harman, it might seem that now I am being unduly demanding. Certainly we have no proof that instructions are used in mediating the organism's behavior. It is entirely possible that what the simulation achieves computationally is achieved in the organism 'non-computationally' (if we assume that such a distinction can be spelled out clearly). But why demand proof? Isn't the mere fact that the simulation uses a set of internally represented rules solid inductive grounds for thinking that the organism does so too? Let me explain why I suspect this line of reasoning. We know that humans can follow instructions; that we can sometimes actively mediate our behavior by applying internally represented rules: when we bake a cake we often follow a recipe, when we type we might use a style-sheet, when we do a proof we sometimes rehearse the rules, etc. But we only do this part-time--it is a moot question (what we're in effect trying to determine here) whether the rest of our behavior is assimilable to these comparatively uncontroversial examples of 'instruction' following. A computer simulation, on the other hand, is a full-time instruction following device. Every output of the simulation is (perhaps tautologically so) the result of the working of its internalized program. This means that simulations of what are patently not instruction-following processes--like Hurricane Han-

nah's swirling--will still look instructional from the simulation's point of view. Meteorologists, as I've said, know to abstract away from the 'instructionality' of their simulation; my claim is that the same sort of abstraction is necessary here.

The case for AC. In this regard, I think that AC is in a better position than AS. AC, we recall, does not focus on what we've termed the 'instructionality' of a simulation, but rather on the specific sorts of operation the instructions instruct. What we still have to ask is how we get from (a) unconscious computational operations going in an organism, to (b) the ascription of tacit knowledge of general computational principles. In my reconstruction of Fodor's argument (AC) I tried to bluff through this last question.

> This (evidence from simulations of depth judgement) gives us a strong inductive basis for saying that the organism being simulated also carries out such computational processes. But we ordinarily explain the ability to perform this sort of complex mental operation by positing some underlying knowledge of general principles. We should therefore use the same explanation here; except that in this case, since the knowledge is not introspectible, it must be tacit.[36]

This formulation contains two questionable steps. The first implicitly identifies computational processes with mental operations; the second makes the assumption that

if a process that mediates an organism's behavior can be described as 'computational' then the organism must be applying general rules or algorithms for carrying out the process. In what follows I hope to show that these assumptions are closely related and both illegitimate.

Tacit knowledge and the knowing subject

The subject and his central nervous system. Whatever philosophers have to say about the issue, ordinary parlance is not very precise about (among other things) where to draw the line between a person and his nervous system. This is not necessarily a fault of our ordinary parlance. We have a large family of fuzzy distinctions which we gainfully employ to mark off these boundaries, but they are not used with much consistency--we're generally agreed to keep them fuzzy. Some popular members of this set are

> (1) the distinction between mental and non-mental events, processes, operations, etc.,
>
> (2) the distinction between cognitive and non-cognitive processes,
>
> (3) the distinction between what a person does and what his body does,
>
> (4) the distinction between how one does something, and what actually goes on when one does that something.

There are no doubt other distinctions that belong on this list, but these will suffice for my purposes.

I've stressed before that to say that a person has some sort of tacit knowledge is to make a claim about his epistemic holdings. Notwithstanding the fuzziness of the notion 'mental', those who are willing to use the term are generally agreed that a state of knowing is a mental state.[37] So if talk of tacit knowledge is to be more than what Nagel calls a 'technical barbarism', it must be shown that in ascribing such knowledge we are saying something relevant about the mental life or make-up of the organism. If we are of a mind to, we can define a sense of 'tacit knowledge' in which a car has tacit knowledge of rules for getting into a higher gear; we could in the same vein add that they have 'memories' of their accidents because dents and smash-ups can be put into one-one correspondence. It is obvious, however, that neither of these ascriptions has any psychological or philosophical significance (except for what it might show about our powers of imagination). What I want to show is that Fodor's analysis of tacit knowledge leads directly to a similar, though less drastic, sort of trivialization.

The problem. Let us return to the homunculus story Fodor uses as an illustration. The real moral of this fable is that unbeknownst to (most of) us, our behavior is 'epistemically' (or: cognitively) controlled from the inside. Although we know that we have no intuitive idea

of how depth estimates are computed, we are asked to imagine that there is a little man (a pseudoperson) who knows a lot about such things and who works for us.

The problem with the argument, figuratively speaking, is that Fodor calls in the little man, has him do his computational/mental business, discharges him, and then wonders whose mental operations are responsible for the behavior. This reverts back to the issue of whether computational processes are ipso facto mental processes (no), but there is a more significant distinction that is being trampled on here: How do we justify the jump from processing that goes on in us, to processes that we perform by applying rules? Where is the argument for redrawing the boundaries of the person somewhere inside the nervous system?

Two arguments. Fodor intimates from the start that an argument from like effects to like causes will provide the mediating link for all this. I want to show here why this is a mistake. Consider Fodor's argument:

> If machines and organisms can produce behaviors of the same type and if descriptions of machine computations in terms of rules, instructions, etc. that they employ are true descriptions of the etiology of their output, then the principle that licenses inferences from like effects to like causes must license us to infer that the tacit knowledge of organisms is represented by the programs of the machines that simulate their behavior.[38]

But this argument will only work if we add the premiss that machines have tacit knowledge of the principles involved in their processing. If this is granted, then the like effects to like causes principle would justify ascriptions to the organism as well. But why should we accept such a premiss? I myself am skeptical of its truth. Although I'm not sure what Fodor would say about this matter, this is certainly not a premiss that can be assumed. On another interpretation of the argument, it begs the question at issue by assuming that the organism has tacit knowledge. What we can say, along with Fodor, is that if we assume that the organism has tacit knowledge, then we can appeal to the simulation for a specification of the content of such knowledge. But we still need a separate argument for making such ascriptions in the first place.

There is one more near-argument that Fodor makes that is worth looking into. In discussing the homunculus, he makes the following claim: "...whenever you X, the little man has access to and employs a manual on X-ing; and surely whatever is his is yours" (Italics mine).[39] If we divide through by the little man, we end up with something like a principle of transitivity: if your psychological faculties can be said to apply some rule r, then the same can be said of you. Here then we

might have the link we've been looking for between processes that go on in the organism and what the organism does. Since Fodor assumes that something like this principle holds, he can pose the key question in the following form: "What are we to say is the epistemic relation an agent necessarily bears to the rules he regularly employs in the integration of his behavior?" (Italics mine).[40] In going from 'A's behavior is integrated on the basis of p' to 'A integrates his behavior by applying p' we are dealing with more than a question of grammatical niceties. What then can be said in support of this principle?

Fodor does not provide any argument in its favor, but I want to consider the following hypothetical line of reasoning:

> There are all these mental (epistemic, cognitive) events that play a part in the production of behavior; mental events have to be part of someone's mental life, they don't just float around unattached; we also know that there is no real little man in our heads. Given all this, what are we to say? Whose mental life could we be talking about if not our own? It must therefore be our secret (read: tacit) mental life.

How to control oxygen levels. This reasoning is flawed, I think, and the way to see that is to consider the following parody of Fodor's position. Consider the metaphor of the army of little men again. We can adapt this model whole to frame an explanation of such phenomena as regu-

lation of oxygen levels in the bloodstream. We designate one little man as representative pro tem for the physiological (including neural) faculties that mediate the integration of the operations involved in regulating oxygen levels (perhaps by controlling breath rate). This little man has a volume entitled How to Control Oxygen Levels which he takes down and uses whenever the organism needs more oxygen in his bloodstream. Given the transitivity principle, we seem driven to the view that the subject tacitly knows the principles of respiratory regulation that he tacitly applies whenever his oxygen levels drop. Although I think common sense revolts against this sort of knowledge ascription, one could bite the bullet and accept this view. Either way, however, Fodor's notion of tacit knowledge loses its interest. If we tacitly know a grammar in the same sense that we know principles for regualting oxygen levels, then there is nothing to alarm the anti-intellectualist.

Notice that this argument about oxygen levels might not apply directly to Fodor's stated criteria which require that tacit knowledge ascriptions only be made to explain (unexplained) know-how. It might be argued that it is somehow improper to say that we 'know how' to regulate our oxygen levels (even though that is what we do when we pant, etc.); that it is merely a physiological process. Given that Fodor provides no guidelines on this

matter, it is not clear that this line of response is open to him. Be that as it may, the issue is beyond that of his explicit criteria. The point is that if we accept the transitivity principle--or something like it--then Fodor's criteria appear to be ad hoc. This is the incongruity we noted earlier between the stated criteria and the arguments offered in their support. Even Fodor must grant that we need the know how/physiological process distinction to determine which phenomena are in the domain of tacit knowledge explanations. How, then, can he ignore it and talk of the organism as integrating his own behavior? Does it make any more sense to say that we integrate our own behavior than that we control our own oxygen levels? To whatever extent a person 'owns' his central nervous system and the computational operations therein, he also owns his respiratory regulation system as well. How will Fodor distinguish the cases?

Relevant and irrelevant distinctions. I want to consider one last response Fodor might make to the effect that the distinction(s) I am trying to protect are in the final analysis not significant for the purposes at hand. This is very close to the position put forward in the following crucial passages from his book The Language of Thought:

> ...while I have argued for a language of thought, what I have really shown is at best

> that there is a language of computation; for thinking is something that organisms do. But the sort of data processes I have been discussing, though they may well go on in the nervous system of organisms, are presumably not, in the most direct sense, attributable to the organisms themselves.
>
> There is, obviously, a horribly difficult problem about what determines what a person (as distinct from his body) did. Many philosophers care terrifically about drawing this distinction, and so they should: it can be crucial in such contexts as the assessment of legal or moral responsibility. It can also be crucial where the goal is phenomenology: i.e. the systematic characterization of the conscious states of the organism. But whatever relevance the distinction between states of the organism and states of its nervous system may have for some purposes, there is no particular reason to suppose that it is relevant to the purposes of cognitive psychology.

> What cognitive psychologists typically try to do is characterize the etiology of behavior in terms of a series of transformations of information. ...roughly speaking, information is said to be available to the organism when the neural event which encodes it is one of the causal determinants of the behavior of the organism. 'Behavior' itself is construed broadly (and intuitively) to include, say, thinking and dreaming but not accelerating when you fall down the stairs.
>
> If one has these ends in view it turns out (again on empirical rather than on conceptual grounds) that the ordinary distinction between what the organism does, knows, thinks, and dreams, and what happens in and to its nervous system does not seem to be frightfully important. The natural kinds, for purposes of theory construction, appear to include some things that the organism does, some things that happen in the nervous system of the organism, and some things that happen in its environment.[41]

Fodor concludes by claiming that "the distinction between actions and happenings isn't a psychological distinction".

I think that much of the foregoing is correct. It is true that the action/happening distinction is not 'psychological' at least in that it is not incumbent on a working psychologist to provide an analysis of these terms; its extension can be treated as intuitively given. It is an entirely different matter, however, to claim that this distinction (and the other related distinctions mentioned above) is "not relevant to the purposes of cognitive psychology".

The example about respiratory regulation should clarify why this is so. What it shows, I think, is that there are <u>two</u> types of underlying processes that Fodor's analysis would run together but which must be kept distinct. Both mediate our functioning in the world, both types involve the operation of the nervous system, and both types can be (perhaps best) understood in terms of computational models specifying successive transformations of information. The first type includes the operations cognitive psychologists tell us mediate our judgements of depth; the second type are like those that physiologists tell us mediate the regulation of our respiratory system. Fodor's criteria can be seen as an attempt to pick out the first set of processes as comprising the tacit knowledge that subtends our know-how. The

second set, on the other hand, dooesn't seem to have much at all to do with our mental/cognitive life. Fodor's criteria should therefore not ascribe tacit knowledge in these cases. The difficulty arises because on Fodor's reasoning, we have no no principled way to distinguish the two sets. How does the cognitive psychologist decide that in the respiratory case we are not dealing with a 'tacit knowledge context' while in the perceptual case we are?

The only available answer I know of is this: we distinguish these processes not at the neurological or computational levels, but in terms of the molar processes they mediate. Processes that fall into the first category subtend the organism's mental life; they are part of the mental sphere and we therefore consider them 'mental'. The processes that belong in the second category subtend functioning that simply goes on in the organism; there is nothing 'mental' about the output of these processes. Unless Fodor is willing to assume some rough characterization of this sort at the molar level, he has no basis for limiting ascriptions of tacit knowledge to cases of know-how. But if he is willing to make the distinction between what we do (estimate depth) and what happens to us (our oxygen levels are monitored) on this surface level, then how can he ignore it at the deeper levels? The point is that although a psychologist

can take these distinctions as intuitively given, it does not follow that they are irrelevant. Quite the opposite; they help delimit the domain of cognitive theory by allowing us to pick out, albeit intuitively, the set of phenomena the theory is to explain.

In conclusion, then, I think that there is a strand in Fodor's thought that can lead to a viable notion of tacit knowledge, but it is not the main strand. The main strand, as I understand it, takes tacit knowledge ascriptions to apply to too wide a domain--the entire range of know how--and leads to the blurring of a number of distinctions that it is in the best interests of psychology to keep sharp--e.g. the distinction between what we do and what happens in us. (To a certain extent, this problem ties up with the simulation approach Fodor selects.) The original goal of using tacit knowledge ascriptions to explain mental competences is to my mind a step in the right direction. Unfortunately, Fodor gives us no hint as to how we are to specify the performances that are to be included under this category: how are we supposed to deal with cases like Harman's cyclist?

We can take a significant step toward answering this question by returning to a passing distinction we made between instructions for doing and principles for deri-

ving. Consider: if you know how to mambo (hardly what we'd want to call a mental competence), then the manifestation of your know-how is normally behavior--mamboing. But if you know how to do a proof in the predicate calculus, then you typically arrive at a proof--we might say a representation of some sort--in a way that you don't arrive at a mambo. As it stands, this distinction is vague and I suspect that it is also open to numerous counterexamples. Nevertheless, it suggests that certain instances of know-how involve what we might loosely call an 'intellectual' output, and it is here that the cyclist and speaker might differ significantly. A distinction close to this one is at the heart of the model for tacit knowledge ascriptions offered by Graves, Katz, et. al. to which we now turn.

NOTES CHAPTER III

1. R. Schwartz, "On Knowing a Grammar" in Hook, ed., Language and Philosophy, pp. 187-188. It should be noted that Schwartz introduces this example to make a slightly different point about physical realizations of rule systems. This point is not duly appreciated by Chomsky in Reflections on Language.

2. This argument is presented quite forcefully in Michael Levin's "Explanation and Prediction in Grammar (and Semantics)", Midwest Studies in Philosophy, v. 2, pp. 135-136.

3. G. Harman, "Psychological Aspects of a Theory of Syntax" in Stephen Stitch, ed., Innate Ideas , pp. 167-168.

4. Harman no longer holds to the specific positions ascribed to him here; nevertheless, the arguments are pertinent and deserve attention. For the shift in Harman's views see his review of Chomsky's Language and Mind in On Noam Chomsky, ed. G. Harman, p. 215ff.

5. For the sake of this argument we will follow Harman in identifying 'system of representation' and 'language'. For contrasting views on the appropriateness of this identification, see Fodor, J. The Language of Thought (NY: Thomas Y. Crowell Co., 1975), and Ronald Arbini, "Comments on Linguistic Competence and Language Acquisition", Synthese, v. 19, nos. 3/4, April 1969, pp. 415-417.

6. Chomsky, "Linguistics and Philosophy" in Hook, ed., Language and Philosophy, p. 88.

7. "Reply to Arbini", p. 427.

8. I do not mean here that the theory is applied in performance--only that it plays a causal role in knowing a language.

9. This hypothetical objector has real-life supporters. See, for example, Goodman's "The Emperor's New Ideas" in Problems and Projects (Indianapolis: The Bobbs-Merrill Co., Inc., 1972), pp. 78-79, S. Stitch's "What Every Speaker Knows", pp. 484ff, and Harman's "Reply to Arbini", p. 426.

10. For a related discussion of this point see Chomsky "Knowledge of Language" in Minnesota Studies in the Phi-

losophy of Science, v. VII, pp. 299-321.

11. Harman, "Psychological Aspects", pp. 172-173.

12. This claim of Harman's was followed by a flurry of countercounter...argument that did not succeed in settling the issue. For a useful overview see M. Atherton, "Tacit Knowledge and Innateness", pp. 3-11.

13. This is quite loose; a state of knowledge is not in itself the cause of anything.

14. See, for instance, Wittgenstein, Blue and Brown Books (NY: Harper Torchbooks, 1958), section I of "The Brown Book" and Philosophical Investigations, I, sections 213-215.

15. "Knowing How and Knowing That" in Collected Papers, v. II (NY: Barnes and Noble, Inc., 1971), pp. 217-218.

16. Ryle is not overly concerned with this problem because of his dispositionalist neo-behaviorist position on knowledge, belief, etc. The problems with this view are well-documented.

17. We might also note that Ryle is wrong in claiming that talk of implicit knowledge must be part of an attempt to reduce know how to know that.

18. Harman, "Psychological Aspects", p. 172. To avoid confusion later: note that Harman adopts a rather idiosyncratic view of internal representation in this paper; see his discussion, pp. 168ff.

19. Chomsky, "Linguistics and Philosophy", p. 87.

20. Bicycle riding has had a strange allure as an analogue/model of linguistic functioning. Consider the following remarks by Jespersen and Russell:

> Do the little brains (of children) think about these different forms and their uses? Or is the learning of language performed as unconsciously as the circulation of the blood or the process of digestion? ...We ride a bicycle without giving a thought to the machine, look around us, talk with a friend, etc., and yet there was a time when every movement had to be mastered by slow and painful efforts, There would be nothing strange in supposing that it is the same with the acquisition of language.

Otto Jespersen, Language: Its Nature, Development, and Origin (NY: W.W. Norton Co., 1964), pp. 180-181,

> ...'meaning' can only be understood if we treat language as a bodily habit, which is learnt just as we learn football or bicycling.

Bertrand Russell, An Outline of Philosophy (NY: Meridian Books, 1960), p. 46.

21. For reasons that elude me (and puzzle Atherton as well (see note 12)), Harman later almost grants that this is so: "I am willing to concede that knowledge of a language is not a form of knowing how in any ordinary sense of this phrase." See his "Linguistic Competence and Empiricism" in Hook Philosophy and Language, p. 145.

22. For a development of this line, see Chomsky's recent shift to 'cognizance' as a replacement for talk of tacit knowledge in Reflections, Part I.

23. Chomsky, "Comments on Harman's Reply", in Hook, Language and Philosophy, pp. 154-155.

24. J. Fodor "The Appeal to Tacit Knowledge in Psychological Explanations", JP, LXV, 20, October 24, 1968, p. 627. Unless otherwise indicated, references to Fodor's views are to those put forward in this paper.

25. Ibid., p. 629.

26. Ibid., p. 630.

27. Optimality, in Fodor's sense, is not a guarantee of truth. A simulation can be optimal in that it specifies only instructions that are elementary for the organism, but the simulated behavior could, in actuality, be a result of a different combination of elementary instructions. The simulation would satisfy Fodor's conditions of optimality, but would fail to tell the true story of how the organism's behavior was brought about.

28. Ibid., p. 627. Fodor goes on to amend this version considerably.

29. See Thomas Nagel "The Boundaries of Inner Space" in JP, LXVI, 14, July 24, 1969, pp. 453-458. This contains a fairly convincing argument that Fodor's notion of complex behavior as a function of elementary operations is not psychologically accurate.

30. "Appeal", p. 632.

31. Ibid., p. 638.

32. The uninteresting exception, of course, is the crackerjack cognitive psychologist--but even here it is

not the explicit knowledge that mediates the behavior.

33. Note that Fodor's criteria are framed in terms of organisms in general, not just humans. This is no doubt an interesting topic--i.e. the tacit knowledge of non-humans--but I will ignore it here. I will use 'organism' only as a liberated (?) way of referring to human beings.

34. I am not sure to what extent traditional behaviorists can accept this sort of talk; the idea is that we 'discover' that the information contained in the stimulus is richer than we had originally supposed.

35. Notice also that AS yields a performance type concept of tacit knowledge that fits very well with the views expressed in Fodor, Fodor Garrett, "The Psychological Unreality of Semantic Representations" in Linguistic Inquiry, Fall 1975, pp. 515-533. AC places tacit knowledge at a one step remove from actual performance; see in this regard J.J. Katz "The Real Status of Semantic Representations" in Linguistic Inquiry, v. 8, number 3, pp. 559-584, and Chomsky's discussion of 'psychological reality' in "On the Biological Basis of Language Capacities" in The Neuropsychology of Language, R.W. Rieber, ed., (NY: Plenum Press, 1976), pp. 4ff.

36. pp. 145-146.

37. For the purposes of this discussion, I have not distinguished between knowledge and belief. Within the context of that distinction it would be more accurate to say that belief is paradigmatically mental.

38. Ibid., p. 640.

39. Ibid., p. 636.

40. Ibid., p. 636.

41. The Language of Thought, p. 52.

CHAPTER IV

THE INFERENTIAL MODEL

Fodor's theory, as we have seen, is beset with difficulties. The crux of the problem, I think, is that he tries to justify ascribing tacit knowledge to a subject on the basis of a computer simulation. If we are already convinced that the subject does have some underlying knowledge, then it is perfectly in order to try to specify that knowledge in terms of a simulation model. But if the question of knowledge is itself moot, there is nothing of direct epistemological value to be learned from a simulation. No matter how we rearrange the details, we cannot get around the fact that simulations, qua simulations, don't really know anything. They only simulate knowers. No one would think that a computer simulation of a 65 MPH car collision causes anyone any injury, and I see no more reason to think that the case is otherwise when we are dealing with a simulation of a speaker. The general point, I think, is that we have to distinguish between a simulation of O and a duplication of O. Given this distinction, Fodor's problem is clear. Unless we can first show that simulations have tacit knowledge, we will be hard put to make the argument from

like effects to like causes work to justify ascription of such knowledge to people.

To avoid misunderstanding, let me add two points. First, I am not arguing that there is some theoretical bar to ascribing full-fledged epistemic states to artifacts of our own creation. The point, rather, is that such ascriptions will be parasitic on the patterns we follow in ascribing such states to each other. What Fodor is asking, if we can oversimplify for a moment, is that we reverse the direction of priority here--i.e. that we ascribe epistemic states to persons on the basis of our discoveries about our artifacts. It is this reversal that I think is unwarranted. Second, I do not claim that there is nothing of import to be learned from a successful computer simulation of a cognitive function. As I have indicated, we can learn that some part of our behavior is (probably) mediated by underlying computational operations, and that (part of) the central nervous system realizes/instantiates a set of general principles for carrying out these operations. It does not follow from this, I have argued, that we know these principles tacitly. If it did--that is, if the proper criterion for ascribing tacit knowledge of p were simply that a simulation of the organism contained p--then we would end up overascribing. We would be driven to the conclusion that, e.g. humans have tacit knowledge of principles

they apply in monitoring the oxygen levels of their blood. Seen in this perspective, the simulation approach turns out to be too liberal.

In this chapter I want to suggest a more conservative model for tacit knowledge ascription suggested by Graves, Katz, et. al.[1] (henceforth 'Graves-Katz') that tries to constrain the range of ascriptions in some principled way. The chapter is divided into three sections. In the first I present the Graves-Katz model, and point out some of the ways in which it differs from Fodor's proposal. In the second, I take up a number of arguments that attack the Graves-Katz model. In the final section I consider and reject two attempts to explain knowledge of language without ascribing any tacit knowledge to the speaker.

The Graves-Katz Model

The role of intuitions

As I indicated earlier, we can view Graves-Katz's proposal as being a direct response to Harman's cyclist argument. Whereas Fodor wants to claim that, contrary to Harman's intuitions, speakers and cyclists both exploit a specific body of tacit knowledge in the exercise of their respective skills, Graves-Katz are willing to accept Harman's intuition. That is, they grant that if the

speaker's knowledge of language were on a par with the cyclist's knowing how to ride a bicycle, then there would be no grounds for ascribing tacit knowledge to either. They argue, however, that there is a difference between the two that has been thus far overlooked.[2] Apart from the speaker's ability to use and understand his language, he also has a set of linguistic intuitions. Speakers are typically able to make judgements about the grammatical properties and relations that hold in their language. They know, for instance, that (1) is grammatical, and that (2) is ambiguous.

(1) *The cat eated the dog

(2) Debbie lives near the bank

The key point to notice here is that such intuitions are clearly instances of explicit propositional knowledge.[3] Furthermore, this knowledge is open-ended; theoretically, there is no upper bound on the number of expressions we can assess in this way.

Tacit knowledge and tacit deduction

Graves-Katz urge that the only way to explain the existence of this explicit linguistic knowledge is to assume that speakers have tacit knowledge of a set of underlying rules from which they tacitly deduce these (explicitly known) results. This leads to the following criterion for tacit knowledge ascriptions:

> ...a proposition is tacitly known if it is one of the propositions appealed to in a tacit deduction; that is, if it is one of the propositions from which the person has deduced a...known proposition.[4]

Cyclists have no unexplained propositional knowledge corresponding to the normal speaker's explicit intuitions. There is therefore no reason to resort to this 'deduction paradigm' to explain anything in the former case. Since Graves-Katz consider tacit knowledge ascriptions only within the framework of this paradigm, Harman's challenge is successfully countered.

All this leaves us with a rather obvious question: Are we really making any conceptual headway in explicating this mysterious notion of tacit knowledge in terms of the equally mysterious notion of a tacit deduction? To go back a step, do we have more reason to believe that people make tacit deductions than that computer simulations have any knowledge? We criticized Fodor in this regard for trying to explain a controversial notion by introducing an explanatory posit that was no less controversial. Isn't the Graves-Katz model guilty of the same thing?

The answer here, I think, is 'no'. The crucial difference is that in Fodor's case, the whole idea that a computer could have any sort of knowledge is moot; this places an extra burden on the inference that it cannot bear. Again, we cannot argue that speaker's must have tacit knowledge of a grammar because of the analogy they bear to simulation models, unless we have reason to say that simulation models have such knowledge. We need more of an argument than Fodor gives to show that we indeed have such reason. On the Graves-Katz model, the speaker's tacit knowledge of a grammar is to be understood on analogy with a bona fide subject of epistemic properties--i.e. a knowing subject who has <u>explicit</u> knowledge of an underlying set of rules; typically, a linguist.

Imagine a field linguist who is attempting to write a grammar of some native language. If he is successful, then he will be able to tell us which strings in the language are ambiguous, which ones are syntactically well-formed, and so on. <u>His</u> explicit linguistic knowledge can be accounted for in terms of the fact that he has explicitly formulated a set of relevant rules, and he (again explicitly) deduces his results from these rules. It is in the comparison of the linguist and the native speaker that Graves-Katz locate the argument from like effects to like causes. They claim that the best expla-

nation of the speaker's explicit linguistic knowledge takes it to be the result of a deduction from this known set of rules. If we accept this postulation, then we have an explanation for the fact that speakers have this set of explicit intuitions about their language. If we reject it, then we have none. As they put it:

> ...the choice before us is whether to extend a familiar paradigm of explanation by postulating an unfamiliar type of knowledge or to save ourselves this postulation by adopting a new paradigm of explanation. We have taken the former alternative because there appears to be no plausible paradigm of explanation available that avoids the postulation of tacit knowledge.[5]

Speakers, Linguists, and Cyclists

This last contention provides us with a convenient way to go at the Graves-Katz model. In this section I want to consider the following charges that might be made against this theory. First, it might be argued that there is really no significant analogy between the native speaker and the linguist, and that Graves-Katz's argument from like effects to like causes is at bottom no better than Fodor's. Second, it could be argued that whatever the status of the speaker-linguist analogy, Graves-Katz fail to establish a significant disanalogy between the speaker and the cyclist (and therefore is again too 'liberal' in the sense specified earlier). In the last

section I take up the view that, contrary to what Graves-Katz assert, there are "plausible paradigm(s) of explanation that avoid the postulation of tacit knowledge."

On the analogy between speaker and linguist

The problem of 'like effects'. Since the 'like effects' in the Graves-Katz model are the speaker's knowledge that (e.g.) a string is grammatical and a linguist's knowledge of the same thing, we have to ask to what extent these effects are really analogous. It is no doubt true that both satisfy the description 'explicitly knows that string s is ambiguous', but this is not enough to make the argument go. Any arbitrarily chosen pair of objects will share certain properties. One could just as well argue that the two are not analogous since only the linguist can tell us something about the principles he used in arriving at his judgement of ambiguity; he, perhaps, could give us the transformational history of the string. The speaker, ex hypothesis, cannot do any of this. How, then, are we to decide if the analogy is significant enough to drive the argument?

We can begin by noting that the notion of a 'like effect' can only be understood relative to a set of theoretical interests and parameters. We cannot expect that every behavioral output available to the linguist will

have its counterpart in the speaker, nor vice versa. But this in itself does not undermine the argument. Our problem then is to demarcate which set of phenomena we will consider the effects, and which we will count as a 'side effects'. This is no simple task--either here or in any other scientific context.

To get some idea as to how we should proceed, consider the case of a rheumatic whose rain predictions mirror the forecasts of professional meteorologists. For the sake of the example, let us also assume that the rheumatic is as mystified by his own accuracy as we are ("I don't know, I just feel like it's gonna rain"). Here we do not want to say that since the meteorologist and the rheumatic evince like effects--they make the same rain predictions--they must therefore have the same causes--e.g. the rheumatic must have access to satellite photographs of the hemisphere, etc., and he makes tacit deductions from this knowledge base. But why not? The obvious reason is that we have other explanations that are less fabulous. Still it should be clear that this tacit knowledge explanation would seem implausible even if we had not yet succeeded in developing a more credible alternative (and there are symmetrical cases where we really haven't--e.g. idiot savants). Two considerations can still be brought to bear here:

(1) If the rheumatic did know (tacitly) what the scientist knows (explicitly) then we could expect the rheumatic to know (explicitly) a great deal more than he in fact does know (explicitly). The hypothetical knowledge base enables the meteorologist to tell us not only whether or not it will rain, but also how heavy it will be, when it will stop, and so on. The rheumatic cannot tell us any of these things. But if they share the same inferential base, then why is the rheumatic's predictive power so poor?

(2) The idea that the rheumatic might come to have knowledge (tacit or explicit) of satellite photos that have never been shown to him is incredible. To take this hypothesis seriously would require a major overhaul in (at least) our conception of knowledge acquisition. Rather than undertake such an overhaul, we look hard for an alternate explanation.

Brentano's solution. The general principles, of which these particular considerations are instances, are included in Brentano's conditions of adequacy taken up in Chapter I.

> (3) ...it is necessary to assume that they (conscious mental phenomena) do not involve at the same time other consequences which are absent at the case at hand, even though there

> is no reason to suspect that they are connected with the concommitant consciousness which is missing in these cases.
>
> (5) ...the origin of the mental phenomena assumed to exist despite the absence of consciousness should not be considered something utterly and entirely inconceivable in itself.[6]

The rheumatic seems to fail the first test because what is "absent at the case at hand"--for instance, knowledge about how heavy the rain will be--seems to be independent of the issue of consciousness. He fails the second test because the origin of his hypothesized underlying knowledge about satellite photos is utterly mysterious.

What I want to do here then is consider the significance of the analogy between the speaker and the linguist in the light of these conditions of adequacy. Graves-Katz's proposal must therefore respond to the following two questions: (i) How do we account for the fact that the linguist's explicit knowledge is so much broader in scope than the speaker's, given that they share the same knowledge base? Can the difference be accounted for in terms of the tacitness of the speaker's underlying knowledge? (ii) We can explain the linguist's acquisition of the necessary underlying knowledge in terms of the general empirical procedures linguist's apply; how do we explain the native speaker's acquisition

of the relevant tacit knowledge? It is to the first of these questions that we now turn.

What every linguist knows, or: why we need linguists

I want to argue that the discrepancy between what the linguist knows and what the speaker knows does not violate Brentano's strictures in the way that the discrepancy in the rheumatic-meteorologist case does. If we take what a competent speaker explicitly knows about the grammatical features of his language and we subtract that from what a linguist knows on the basis of an adequate grammar that he has devised, we can expect the remainder to fall into two sets. First, there can be a set of grammatical judgements (possibly empty?) which the linguist can make that the native cannot. Second, we can expect there to be a class of facts/propositions that are available to the linguist but not to the native speaker. As an example of the first, imagine that we ask both subjects if a multiply center-embedded sentence is equivalent in meaning to its unembedded counterpart. The linguist might be able to figure this out, while the native might not. As an example of the second, we can imagine that the linguist could determine not only that some string *s* is ambiguous, but that it involves *n* applications of the dative movement rule in its syntactic derivation. This sort of knowledge is not available to

the ordinary speaker. To support the argument from analogy, we must somehow explain away these (seeming) discrepancies.

As far as the first set goes, it seems plausible to think that the shortcomings on the speaker's part are explainable in terms of processing difficulties having to do with memory limitations and the like. Since the linguist has the grammar in front of him, he can overcome these limitations and arrive at a judgement. If this is in fact a step toward the correct explanation, then this aspect of the linguist's 'extra knowledge' poses no real problem for the Graves-Katz model. The deficit could be chalked up to what Brentano refers to as the absence (in the case of the speaker) of the concommitant consciousness that is present in the linguist's case.

In order to see why the second set poses no real problem either we have to attend to the distinction between the deductive consequences derivable from/in a theory, the propositions of the theory, and the metatheoretical truths about the theory. To see the relevance of this distinction, consider the following question (which might have occurred to the reader): If every speaker has tacit knowledge of a grammar, and every speaker can make tacit deductions from what he tacitly knows, then how come we don't come to have explicit

knowledge of the grammar on the basis of our tacit deductions? From our tacit knowledge of G we ought to be able to deduce G and thereby arrive at explicit knowledge of G (by the rule p/therefore p). Why then do we need linguists?

The distinction I have in mind would help to explain this oddity. Let us assume for a moment that we know some natural deduction system for first-order logic and that we can apply it perfectly. The deductive consequences derivable from the theory can be represented as ordered pairs whose first member is an argument, and whose second member is the value 'valid' or 'invalid'. If we happen to have the theory written out in front of us, then we can not only decide that a certain pair is generated by the theory, but we can also recite the rules of implication that the theory allows. Along the same lines, we can determine that there are so many applications of rule r in this proof, that the theory is complete, that it is sound, and so on. The latter are clearly metatheoretical results, only the ordered pair of argument and value is a deductive consequences in the theory. If we built a logic machine to solve problems having to do with validity, we might have no cause to include the facility to produce any sort of metatheoretical statement as output. Similarly, if we decide that human problem-solvers develop an internalized natural

deduction system, we need not expect that any metatheoretical results would be explicit.

The carryover to the case of tacit knowledge of a grammar should be obvious. In ascribing such knowledge we are only committed to the view that speakers can make tacit deductions within a grammar; that is, they can make inferences to the deductive consequences of the theory. (Strictly speaking it is not the grammar itself that yields the deductive consequences that we've been considering--e.g. that (1) is ungrammatical-- but rather the grammar plus the linguistic theory; the latter contains the definitions of terms like 'ambiguous', 'syntactically well formed', etc.) The fact that a string can be derived by way of a dative movement rule is no more a deductive consequence of the theory (a grammar) than the fact that an argument can be shown to be valid by the successive application of the modus ponens rule is a theorem of first order logic. Both results are metatheoretical, and since metatheoretical reasoning depends on the availability of the theory for scanning, we have every reason to expect that the linguist, but not the native speaker, will have this sort of knowledge. This difference, then, can be accounted for on the basis of the fact that the speaker's knowledge is tacit.

In order to satisfy Brentano's condition and thereby make good the contention that the analogy between the speaker and the linguist is significant enough to warrant ascription of knowledge of a grammar to the former on the basis of the 'observable' knowledge of the latter, we still have to deal with the problem of acquisition--point (ii). I want to delay discussion of this issue, however, since the problem of acquisition will reemerge more forcefully when we treat alternatives to tacit knowledge explanations in the latter part of this chapter. Now I want to turn to a different line of attack.

On the disanalogy between speaker and cyclist

One way to view Graves-Katz's inferential model is to see it as a claim to the effect that if we want to take knowledge of language to be a form of knowing-how, then we have to admit that it is a very special sort of know-how. In other words, that there are forms of know-how, of which knowledge of language is one, that involve a critical admixture of knowing-that. This then is what distinguishes cycling from linguistic functioning--the former has no relevant propositional knowledge. What I want to do in this section is introduce a number of arguments that reject this view. All are designed to show that in the end there is really nothing about the speaker that makes it necessary for us to ascribe tacit

knowledge to him but withhold it from the cyclist. First I consider the claim that the explicit intuitions of the speaker do have their counterpart in the cyclist's case as well. If this is so, then the inferential model would force us to ascribe tacit knowledge to cyclists as well. I then go on to consider the other side of the coin--i.e. two types of skepticism about the claim that speakers really have any special explicit linguistic knowledge. The first rejects the idea that speakers know anything about grammatical properties of expressions. It is suggested that what they really know is only that certain strings are acceptable, similar to other strings, and so on. It is then claimed that this sort of explicit knowledge is explainable without recourse to tacit knowledge of a grammar. The second skeptical argument focuses on the claim that what speakers have is knowledge. They might very well believe that certain strings are grammatical, ambiguous, etc., but the justifications for these claims are typically not available. Here we will have to see whether this argument poses any real problem for the inferential account suggested by Graves-Katz.

What every cyclist knows. Every cyclist knows, for instance, that

> (3) if he pedals backward his bicycle will not go forward
>
> (4) his bicycle will move in the direction he turns the handlebars

(5) if his front wheel comes off he will crash.

Consider the following claim: How is it, given incontestable instances of explicit propositional knowledge like (3)-(5), that Graves-Katz can claim that in the cyclist's case there is no explicit propositional knowledge to be explained?[7] Of course, what the cyclist knows is not as complex as what the speaker knows, but this would seem to be a matter of the relative complexity of language versus cycling. But why should we say that speakers have a special sort of knowledge? The difference between knowing how to ride a bicycle and mastering a language, according to this counterargument, is simply a matter of degree.

Even though I think that the premisses of this argument are all true, it nevertheless falls short because it fails to recognize a radical difference in kind between the sort of knowledge we want to attribute to speakers and the sort of knowledge we attribute to bicycle riders. Specifically, we can say that speakers have two different kinds of explicit propositional knowledge relating to their language. Aside from their knowledge that (1) is ambiguous and (2) ungrammatical, speakers of English usually also know that something like (6) is true.

(6) Most pluralized nouns (in English) end in 's'.

(6), we should note, is very different from (1) and (2). In the first place, it expresses an inductive generalization; if we turn it into a categorical--replace 'most' with 'all'--it becomes false. As is the case with all rules of thumb of this sort, rigidly consistent application will lead to error: consider 'sheep'.[8]

The judgement that (1) is ungrammatical, on the other hand, is categorical. Notice also that use of (1) in normal conversation involves violation of a rule of English, and is therefore subject to correction. This again, is not the case with (6). If we correct someone for saying 'radiuses' instead of 'radii' we are not censuring them for violating (6), but rather for violating some specific rule for forming the plural of words like 'radius'. It makes no sense to talk of violations of (6)--it's not that kind of rule.

One moral to be drawn is this. In all cases of know-how, it is likely that a reflective agent will come to notice certain patterns that emerge in the exercise of his skill. A chess player might discover that the best strategy for winning is control of the center, for example, and a bridge player should eventually learn that he would do best not to open unless he has an available rebid. These generalizations are usually based on evi-

dence acquired in the practice of the skill, and are important for continued success in practice. Since a large part of mastery of language involves the acquisition and development of communicative skills, making the appropriate generalizations of this sort can be important to speakers as well. This sort of explicit propositional knowledge, however, is parasitic on the 'core' know-how the speaker already possesses. We would certainly be surprised if such knowledge was consistently absent in any case--e.g. if an English speaker had to think a while to determine how most nouns in English are pluralized--but as long as the person pluralized nouns correctly we would class him as fluent. Contrast this with the knowledge of the grammatical properties of (1) and (2). If someone can't tell the difference between a grammatical sentence and an ungrammatical one--e.g. if (1) seems OK to him--then in an important sense he has not mastered the language.[9]

The counterargument fails, then, because it assumes that the only type of explicit linguistic knowledge speakers have is on the order of (6). But this is a mistake. So even though the cyclist does have explicit knowledge on a par with the speaker's knowledge of facts like (6), it is not this knowledge that prompts the tacit knowledge hypothesis. What distinguishes speakers from cyclists are intuitions like the one that (e.g.) (1) is

ungrammatical.

On skepticism with regard to explicit grammatical knowledge. For Graves-Katz, ascriptions of tacit knowledge are necessary because speakers exhibit a level of explicit knowledge about the grammatical properties of (bits of) their language that cannot be explained otherwise. Even if we put aside the philosophical problem about tacit knowledge for a moment, we should notice that an adequate characterization of what it is that a speaker knows about his language has important methodological ramifications for the working linguist as well. Intuitions about the grammatical properties of strings are regularly employed in testing theoretical hypotheses. In order to use these intuitions as empirical input to the task of grammar/theory construction, we have to have a clear sense of what it is that the informant tells the researcher when he offers such a judgement.

In response to Graves-Katz, it has been claimed that speakers don't really have any explicit knowledge about grammatical properties of their language. In this section, then, I want to examine some alternative views as what it is that speakers, qua speakers, know about their language:

> (i) Speakers only know that certain strings are acceptable in ordinary discourse while others are not.
>
> (ii) Speakers know only that certain strings are similar to others.

(iii) Speakers know that strings have specific grammatical properties--e.g. (un)grammatical, (un)ambiguous, et. al.

The first position has it that speakers cannot really make judgements of grammaticality at all; the second, that even if such judgements are possible, they can only be made indirectly; and the third (Graves-Katz) that such judgements can be made pretty much directly. I should point out that this list is by no means exhaustive. One could probably formulate sharper distinctions between (e.g.) judging that a string is not well-formed, that it is ungrammatical, that it is odd, deviant, bizarre, and so on. My main interest here is in the contrast between the first two alternatives and the third.

Grammars and grammaticality. On the face of it, it seems odd to deny that speakers know that strings like (1) are ungrammatical. I certainly would find it odd if my claim to know this was challenged. The philosophical import of this point is frequently overlooked; I would argue that the fact that my claim is generally accepted provides at least a prima facie reason for thinking that I in fact <u>do</u> know what I claim to know. The burden of proof rests on the skeptic; not on me. Perhaps a part of the reason this point is easily missed is that one can

be overly taken with the philosophical quagmires attending the problem of explicating 'A knows that p'. But it would be a mistake to think that these problems of explication cast a shadow on all our knowledge claims. In fact, quite the opposite is true. In order to even begin to construct some theory that explicates 'A knows that p', there must be a pretheoretically clear set of cases where such knowledge claims are considered appropriate. Until we have reason to think otherwise, we ought to include my claim to know that (1) is grammatical in this set.

The skeptic, though, might contend that there is reason to think otherwise. Consider the following thought-experiment. I am given a list of strings, and I am asked to mark them either 'grammatical' or 'ungrammatical'. Assume that (1) is on this list and that I mark it 'ungrammatical'. Now it could turn out that the best grammar that the best linguist can formulate generates all and only the strings I have marked grammatical, except for (1). Here it runs counter to my judgements, and generates (1) as grammatical. Since proferring judgements is one type of linguistic performance, we can imagine that the best performance theory available offers a viable explanation as to why I (mistakenly) failed to mark (1) as grammatical.[10] The completed 'best theory', then, implies that I

misclassified (1), but gives a non-ad hoc explanation of my misclassification. Now we are faced with a choice. We can, on the one hand, reject this theory on the grounds that it fails to match my prediction of (1) as ungrammatical. But this attitude would be tantamount to treating my classifications as some sort of eidetic intuitions. This, I think, is unwarranted. A theory of visual perception, for instance, need not make false predictions so as to duplicate my errors in judgement in cases of (e.g.) optical illusions; it only needs to explain why I make these errors. We should expect no more from a grammar-plus-performance theory. On the other hand, if we take the theory to be correct as it stands, we have to give up the view that I know that (1) is grammatical. I cannot know what isn't so.[11] But now the skeptic steps in and asks us to generalize this last point. If our knowledge claims in this area must wait on theoretical confirmation, we can justify our claims until we can invoke the support of these 'best theories'. In the light of this, should the intuitive grammatical knowledge claims of ordinary speakers be taken at face value?

Steven Stitch, for one, thinks that they should not. The moral he would draw from our thought-experiment is this:

> Grammaticality, ambiguity, being a noun phrase and so forth, are theoretical properties of expressions. ...Expressions in a language have these properties and share these relations if they are said to have them by the correct grammar of the language. To be grammatical is to be classified as grammatical by a correct grammar. And a grammar is the correct grammar of a language if it is the best grammar that accounts for all the data and meets certain further constraints imposed by a general theory of grammar. But if this is correct, it seems odd to say that speakers judge expressions to be grammatical sentences, subject of a given sentence, related as active and passive, and so forth. For these judgements could be reasonably justified only after painstaking construction of a grammar. So it would appear that if the judgements that are the grammarian's data are judgements about grammaticality and other grammatical properties and relations, they are judgements that demand justification not available to the ordinary speaker.[12]

A similar point, we may note, is made by Putnam:

> Presumably, then, to 'intuit' (or assert, or conjecture, etc.) that a sentence is grammatical is to 'intuit' (or assert, or conjecture, etc.) that the sentence is generated by the highest valued Gi in the class Σ which is such that it generates all the grammatical sentence types with which we have been supplied by the 'input' and none of the ungrammatical sentence types listed in the 'input'.[13]

Stitch's point seems to be that judging an expression grammatical can only be the business of theoreticians--in this case linguists--and that speakers must therefore be doing something else. The fact that the speaker uses the term 'grammatical' when he responds to the linguist's query should not be taken to imply that he is sensitive to the real grammatical status of expressions. What we

have is (property-) reference under a false description.

The weak spot in this argument, as I see it, is Stitch's notion of grammaticality as a theoretical property. To see why this is so, think of the properties associated with 'is prime' (as applied to numbers) and 'weighs more than'. There is a sense in which we can say that these too are theoretical properties that can be satisfied by numbers and ordered pairs of physical objects. Number theory and physics are no doubt theories, and these properties are defined within those theories. Does it follow that the typical man on the street--no matter what he says-- doesn't really know that three is prime? It seems that by parity of reasoning Stitch would have to say 'yes'. "To be prime", he should say, "is to be classified as prime by the best number theory." Since most people know hardly any number theory we would have to somehow reinterpret what they say: either they don't really know what they claim to know or they are using 'prime' in some non-standard way. The situation is stranger still in the case of a property like 'weighs more than'. Couldn't one know that one thing is heavier than another without having any conception of a 'best physical theory'?

There is certainly a close relation between being grammatical and being classed as such by a correct grammar, but the relation is not what Stitch takes it to be. We could just as well say that to be grammatical is to have the structure of a sentence the MLA style manual (or even Edwin Newman, for that matter) suggests we use in writing papers. It would clearly not do to argue that most people (and this might even include many linguists) cannot judge grammaticality because they are unfamiliar with the MLA style manual and its suggestions. (After all, the writers of the manual have to be constrained in some way too. Perhaps they consult with Mr. Newman.) I think the relation between grammaticality and a grammar can be better understood along the following lines.

> To be grammatical in L is to have the structure of a sentence of L. The definition "x is grammatical in L=df. an optimal grammar of L generates x" gives an explication of the ordinary notion of grammaticality. It does not eliminate the ordinary notion.[14]

A fuller treatment of these issues would have to clarify the distinction between an explication, a definition, and so on, but that is a separate matter. The relevant point here is that in the sense in which grammaticality <u>is</u> a a theoretical property, it does not follow that one cannot make judgements about it unless one knows the theory.

What the thought-experiment does indicate, I think, is that our explicit linguistic intuitions are not incorrigible: that is, in certain sorts of cases they can be overruled. But this is not the same as saying that such intuitions cannot be instances of knowledge and that we have no knowledge even in the other cases. The fact that I mistakenly claim (over and over) that 'sacrilegious' is spelled s-a-c-r-e-l-e-g-i-o-u-s does not in any way jeapordize my knowing that 'cat' is spelled c-a-t. It's not that spelling and grammar are substantively similar. The point is more general--we can know that something satisfies a certain predicate ('has a certain property', if you like) even though we sometimes misapply that predicate.

Grammaticality and acceptability. It is a fairly common assumption among language researchers that we cannot legitimately expect native speakers to provide grammaticality; that we can elicit judgements of acceptability and no more. One possible reason for this assumption might be the (undeserved) credence given to the claims about grammaticality-as-theoretical-property that were rejected in the last section. No doubt there are other ideological reasons as well. I think that to some extent it is probably true that unless given very specific instructions, speakers will tend to interpret

questions about grammaticality as questions about acceptability. What I want to argue, however, is that the speaker's ability to make acceptability judgements does not imply anything one way or the other about his intuitions about grammaticality.

Again, a comprehensive discussion of these methodological issues would take us beyond the scope of our concerns, so I will just make one small point. As before, the burden of proof remains on the shoulder's of the skeptic. It is a matter of empirical fact that in a great number of instances speakers do make judgements of grammaticality that are independent of matters of acceptability. The reader will certainly find that he too can make such judgements for an indefinite number of cases. Most speakers will even tell us that in certain sorts of situations they will accept ungrammatical strings and in others find certain grammatical strings acceptable. There should be nothing surprising about any of this.

As an aside, we might note that the situation here is reminiscent of what we find when we ask subjects to pass judgement on the morality of a certain action in a given situation. Much to the chagrin of moral philosophers, subjects often decide on the 'acceptability' of an action independently of their assessment of its moral

status ("Well, that's still OK even though its morally wrong.") One often hears the claim (from students) that although its wrong to keep the money when a cashier makes a mistake in your favor, its nevertheless acceptable. (The bigger the concern, the more acceptable; the 'wrongness', of course, remains pretty much constant.) There is an almost irresistible temptation to claim that these students have not yet learned the meaning of the moral terms 'right', 'wrong', etc., but it is not even clear how we would go about showing this.

The upshot of our discussion thus far is that the speaker's intuitions about the grammatical properties of expressions of his language have at least a prima facie claim to being knowledge, and we have as yet found no reason to reject that claim. What we still have to consider is the idea that speakers can only make judgements of similarity; this, in effect, is what Stitch ultimately contends:

> Speakers are able to recognize strings as similar or different. Given a few examples they know how to classify others. And knowing how is quite generally associated with knowing that, provided we are careful to describe accurately what is known. The Bedouin may not know that you turn 30 degrees NNW. He does know that you turn <u>this way</u>. The speaker need not know that a string is grammatical. But it seems quite plausible to say he knows that this string is like those others.[15]

I want to defer discussion of this point till the next section in which we will take up Stitch's specific explanation of how speaker's come to have this knowledge of similarity relations between sentences. Before turning to this issue I want to make a few remarks about grammatical intuitions vis-a-vis the distinction between knowledge and belief.

Explicit knowledge and explicit belief. So far we have been concerned with the appropriate specification of the content of a speaker's explicit linguistic knowledge. As I indicated earlier, however, it might be argued that speakers have no explicit propositional knowledge of this sort at all. Intuitions of grammatical properties, on this view, only have the status of explicit beliefs. 'Intuiting' that an expression is ambiguous is not the same as knowing it to be so; what is missing is justification. This too is an issue that will be taken up later, so here I will limit myself to three small points.

In the first place, what was said earlier bears reiteration. In order to provide a context for the analysis of knowledge we have to accept a set of clear cases where knowledge ascriptions are paradigmatically in place. In ordinary discourse, my claim to know that (1)

is ungrammatical is a non-controversial member of this set.

Secondly, it must still be shown that the distinction between explicit knowledge and explicit belief has any real bearing on the argument under consideration. Recall, that what we are after at this point is a significant disanalogy between speakers and cyclists that could support the ascription of underlying tacit knowledge to explain the performance of the former but not the latter. We as yet have no reason to think that treating grammatical intuitions as explicit beliefs will in any way undermine the argument. Unless it can be shown that cyclists have a body of explicit beliefs that cannot be explained unless we hypothesize a level of tacit knowledge of (something like) the principles of mechanics, then the belief/knowledge distinction does not seem to me to be especially relevant here.

Finally I want to mention a novel argument that purports to show that speakers must have explicit knowledge of (at least) some of the grammatical properties of expressions in their language. This argument is presented by Kitcher:

> ...one can claim that linguistic knowledge is intricately involved in many other types of knowledge. To take an obvious example, our knowledge that p may depend on our knowing that an argument for q is fallacious, and this latter knowledge may depend on our knowing that a particular expression is ambiguous or belongs to a particular grammatical category.[16]

To take a concrete example, we might know that Jimmy Carter is the president, and that the president gets elected every four years. We do not draw the conclusion that Jimmy Carter gets elected every four years. Our knowledge that this conclusion is not true depends on our knowing that the expression 'the president' falls into one grammatical category in the first clause and a different category in the second. So if we know that the argument is invalid, then we must also know this grammatical fact. This argument assumes a great deal about the relation of logic and grammar, and I leave it to the reader to decide whether it is sound. We are now ready to turn to the question of alternative explanations.

Alternative Explanations

Even if we grant that knowing a language involves a level of explicit grammatical knowledge, we have still not made the case for ascribing tacit knowledge of a grammar. Besides showing that such an ascription would explain the fact that people have such knowledge, we must also show that there are no simpler explanations forthcoming. As I indicated, this too is one of Brentano's conditions of adequacy on postulations of unconscious knowledge--i.e. that

> (6) ...the phenomena under discussion cannot be understood, at least not without the greatest improbability, on the basis of other hypotheses.[17]

In what remains of this chapter I will be looking at two closely related alternative hypotheses. I begin with a perceptual knowledge model suggested by Stitch, and go on from there to discuss Kitcher's attempt to provide an updated dispositionalist account.

The Perceptual Model

Stitch's account. I mentioned earlier that according to Stitch the grammatical knowledge of speakers is ultimately a matter of knowing that "this string is like those others." In keeping with this characterization it is natural that he should try to explain speakers' intuitions along the lines of perceptual knowledge:

> Much as a person may come to know something about the color of an object by seeing it, or about the taste of a liquid by tasting it, so I maintain that he may come to know about the properties and relations of sentences by hearing the sentences.[18]

On this view, learning a language consists in the acquisition of

> the perceptual or cognitive capacity to acquire propositional knowledge in certain circumstances. ...grammar is the study of one such capacity and is thus a branch of psychology substantially similar to the study of other sorts of perceptual and cognitive discriminatory abilities.[19]

What we acquire when we learn a language is the perceptual ability to to classify strings in certain ways. Just as we can make similarity judgements about the shapes and colors of physical objects, we can sort linguistic expressions in terms of the grammatico-perceptual properties we perceive them to have.

As Graves-Katz point out, it is hard to tell whether Stitch thinks that linguistic intuitions are part of our perceptual knowledge of the world, or only that they can be understood on analogy with such knowledge. In either case, his point is that knowledge of this sort can be accounted for without introducing the notion of underlying tacit knowledge from which it is deduced. We know that (1) is ungrammatical, and we know it even though we can present no evidence to support our claim:

> As Chisholm would have it, what justifies me in counting it as evident that I seem to see a yellow object (or that I have a dull pain) is simply the fact that I seem to see a yellow object (or have a dull pain).[20]

The same sort of explanation would apply here as well: the explicit linguistic knowledge we have is 'directly evident'. The speaker is exposed to an expression like (1) and 'just sees', as it were, that the expression is similar to the members of some set s.

Problems for Stitch. As I see it, there are two critical premisses supporting Stitch's argument, and they are both of questionable cogency.

> (1) that linguistic intuitions are a species of perceptual knowledge, or at least that there is a significant analogy between the two, and
>
> (2) that perceptual judgements are directly evident and therefore not be based on inferences from what the subject tacitly knows.

Let me explain my misgivings.

As far as the first premiss is concerned, I think that Stitch is partially correct in that there is a connection between linguistic intuition and perceptual knowledge. But the connection is not what he makes it out to be. We can begin to get clear on what the connection is by considering one very significant disanalogy between the two sorts of cases. Imagine that we are given a bunch of things and that we are asked to tell an experimenter what color each one is. When we report that the first item is blue, we are normally conveying something that 'we just found out'. This seems to be a general truth about perceptual knowledge; perception is so wonderful because it usually carries us from a state of ignorance to a state of knowledge. If Stitch is right, then the same ought to hold true of the

knowledge expressed in linguistic intuitions. It is obvious, though, that this is not the case. If I present a speaker with the string:

(0) *Be peace no there

and proceed to elicit his judgement to the effect that it is ungrammatical, has he found out something he didn't know a moment ago? I would maintain that even if he has never seen the string before, the subject will deny that he has just now discovered that (0) is ungrammatical. He is more likely to say that he knew it all along, and I think that in the relevant sense he would be right.

Stitch rejects this possibility because he is under the impression that the only alternative to the perceptual model is the idea that when we question the subject about the grammatical properties we are simply tapping into long-standing beliefs (often misleadingly called 'dispositional') and making them 'occurrent' in consciousness. We either admit that the knowledge is of 'recent vintage' or else we have to claim that it is 'long-standing'. My view is that both these alternatives are misleading, and that both obscure the relationship between perceptual knowledge and linguistic intuition. Let me explain why.

There is no doubt that what the subject does when he provides us with linguistic intuitions involves perception. Simply put, perceptual mechanisms are required to

alert him to the fact that the string we are presenting is a token of a certain type. "That this concatenation of sounds is an instance of the string 'Be peace no there'", we might say, is something he 'just found out'. But it does not follow from this that the speaker's knowledge about the nongrammaticality of (0) is of recent vintage.

Two rather commonplace distinctions will help to clarify this matter. I've already used this distinction between knowledge of the properties of sentence tokens (loudness, neatness of print, color on ink, etc.) and knowledge of the properties of sentence types (grammatical, imperative, overused, etc.). (A more complete taxonomy would have to include properties of types/tokens in context as well.) In judging (0) to be ungrammatical, the speaker is judging it only qua token of a specific type. In making this judgement he is committed to the claim that every set of ink-marks, sound waves, or whatever, that exemplifies this sentence type is grammatical. Linguistics would make very little sense otherwise. Stitch misses this crucial point. For him, grammatical intuitions are things like "one of a pair of sentences sounds more similar to a given sentence than does the other".[21] But sounding alike is not what we are interested in when we ask about properties of types (things are not as clear when we deal with phonology, of course).

It's no wonder then that Stitch leans toward according intuitions the status of incorrigible judgements.

But now it seems that we've painted ourselves into a corner with these distinctions. We surely want to grant that the knowledge expressed in the proferred intuition about (0) was not in the forefront of the subject's consciousness before we prompted him with the string. Everyone accepts this fact. At the same time we don't want to say (i) that the knowledge is of recent vintage in the sense that prompting allows the subject to discover something new--i.e. that '*Be peace no there' is ungrammatical, or (ii) that it is long standing in the sense that, to oversimplify, it was simply sitting there in long term memory waiting to be summoned (like my knowledge that Helena is the capital of Montana). What then are we to make of linguistic intuitions?

It is here that we have to introduce a second distinction between knowledge that is 'latent' and knowledge that is 'non-latent' (I adopt this thoroughly inelegant terminology because I know of no other dichotomous pair that is at once both more elegant and less misleading. Consider: occurrent/dispositional, actual/potential, explicit/tacit (already taken), performance/competence). This distinction is the one that explains the fact that we all know that dogs don't come from the Dog Star even

though only a few of us have bothered to give this matter any occurrent/explicit/actual thought. It is really no news to us if someone informs us that canines do not have this exotic origin. I would claim that if I had died never having had any thought relating the Dog Star to dogs, I still would know this.

There really is nothing very mysterious about this. We all have a large store of general facts at our disposal--our 'belief network'-- and this network rules out the possibility (latently) that dogs are anything but what they seem to be. What is still a problem, and what I am not able to solve here, is how far we can stretch our network by inference. There are certainly things that follow from what we know that we don't know (unless we've learned them non-latently). For instance, one does not know all the truths of arithmetic in virtue of the fact that one is familiar with Peano's axioms. On the other hand, one knows about the non-connection between dogs and the Dog Star even though it might not follow deductively from what one does know. Generally speaking, it seems that knowledge is not closed under induction or deduction. Nevertheless, it certainly covers a broader range than the set of propositions we've consciously considered and accepted as true. That we don't know where to draw the line does not imply that there is no line to be drawn.

We can sum up our discussion of Stitch's first premise by saying that although linguistic intuitions involve perceptual judgements, perception is no more an appropriate model for linguistic knowledge than it is for moral knowledge. In judging whether a given act is moral I also have to hear, see, and so on, to pick out the relevant features. Be that as it may, moral judgements--whatever they are-- are not perceptual judgements.

As for Stitch's second key assumption, it seems to me that it rests on a identification of 'directly evident' and 'known directly'. As I understand the terms, the first is most appropriate where we have no awareness of any deliberative process in arriving at a judgement that something is the case. A great many perceptual judgements have this property, but not all: 'I seem to see a yellow object' is not the same in this regard as 'I seem to see a banana-colored object' (where we might have to squint). Colloquially, if you have to think about it, then it isn't directly evident. Direct knowledge seems to me to be an entirely different matter. To say that a bit of knowledge is direct seems to connote that it was not arrived at by using any inferential process. Whether the two notions actually coincide seems to be an empirical question.

As things stand now, one would have to say that the facts tell against Stitch's position. A great deal of current theorizing in cognitive psychology makes it look likely that whether or not a given claim is directly evident for someone, it might still involve an underlying inferential process. As we saw in our discussion of Fodor, this hypothesis is a dominant approach to research in perception as well. Even if the subject tells us that it is directly evident to him that he seems to see a yellow patch, it does not follow that his knowledge cannot be inferentially derived from some underlying tacit knowledge. So even if Stitch were right in assimilating linguistic intuitions to perceptual knowledge, this would not yet count as a genuine alternative to Graves-Katz's model of explanation. We now turn to a second alternative.

A dispositionalist alternative

In this section I take up Kitcher's attempt to explain the speaker's explicit linguistic knowledge without introducing the notion of tacit knowledge. In the first part I present Kitcher's argument against the position defended by Graves-Katz. This argument touches on the question of acquisition that was deferred earlier. In the next section I present Kitcher's revised

dispositionalist account of knowledge of language, and I discuss its relation to the causal theory of knowledge. In the final section I present a number of criticisms of Kitcher's theory.

Kitcher's argument. Although our focus has been almost exclusively on Graves-Katz's model for tacit knowledge ascriptions, such ascriptions are for them but an intermediate step in a more important project--i.e. the defense of nativism.[22]

> In order for rationalists to argue from the existence of innate structures to the existence of innate knowledge, some notion of tacit knowledge is required so that they can claim that the grammatical structures represented in transformational grammars are objects of knowledge.[23]

If this last step is granted, then it can be argued that the only way speakers could come to have tacit knowledge of a grammar of a specific language is if they have some prior tacit knowledge of a universal grammar which makes grammar-acquisition possible. This universal grammar (often misleadingly called 'linguistic theory') is what nativists want to claim is innate. So the argument ultimately has two steps:

(1) from explicit linguistic intuitions to tacit knowledge of an acquired specific grammar, and

(2) from tacit knowledge of an acquired specific grammar to tacit knowledge of an innate universal grammar.

Step (2), we might note, can be considered a skeleton response to Brentano's condition of adequacy to the effect that the origin of the unconscious psychological phenomena under consideration should not be "utterly and entirely inconceivable in itself".

What Kitcher wants to argue is that (1) and (2) don't mix. That is to say, if we are won over by the argument offered in support of (1), then we ought to reject the argument for (2); conversely, if we take the argument for (2) to be valid, then we are compelled to reject the argument for (1) and thereby deny the soundness of (2):

> Here, then, is the nativist's dilemma. Either his argument for the attribution of tacit knowledge is correct but we have no innate knowledge; or that argument is flawed but we are not compelled to accept the attribution of tacit knowledge to mature speakers.[24]

Kitcher's underlying reasoning is something like this. What is special about the Graves-Katz model is that it directs our attention to the one aspect of knowledge of language that the 'ability approach' could not handle. The ability approach, which we have discussed in the context of Harman's earlier views and Ryle's position, tries to treat knowledge of language as a set of inter-

locking dispositions to respond in certain ways to stimuli. What this account misses, however, is the fact that an informant's intuitions about his language are more than responses to stimuli; they count as knowledge. In order to explain their status as knowledge one needs more than a causal account of how the relevant belief that p was produced-- it is also necessary to show that the person is justified in believing p.[25]

Here there seem to be two ways to satisfy this justification requirement. Either we show that the belief is justified because it is properly inferred from some other justified belief(s), or we claim that the belief is directly justified. Graves-Katz, of course, argue that our explicit linguistic beliefs count as knowledge because they are tacitly inferred from some prior justified beliefs which include the specific grammar of the language. If we ask how this specific grammar satisfies the justification requirement, the answer again is that it too is derived on the basis of a tacit deduction from a set containing the universal grammar. But what if we now apply the same demands on the hypothesis that speakers have tacit knowledge of a universal grammar (the heart of the nativist claim)? It is here that Kitcher locates the dilemma. We certainly have no more rungs left on the ladder of <u>in</u>direct inferential justification. But by the same token, the claim that

universal grammar is directly justified does not seem plausible either: "The classical articulation of the analysis of knowledge as justified true belief allows for direct justification of perceptual beliefs: my seeing yellow justifies directly my belief that I see yellow. A complex belief, such as the belief that the grammars of all natural languages have a phrase-structure part, seems very different from such examples."[26] If, notwithstanding this 'classical articulation, we decide to bite the bullet and accept universal grammar as directly evident, we might just as well take the speaker's linguistic intuitions to be directly evident and never begin the ascent up the ladder.

The argument, in short, is this. Tacit knowledge of universal grammar is only brought in by Graves-Katz to explain tacit knowledge of a specific grammar. But if the argument for tacit knowledge of universal grammar is correct, then there is no need to talk about tacit knowledge of a specific grammar at all. To adapt Wittgenstein's remark, the nativist throws away his ladder before he finishes his climb.

Once the dust clears, we can see that we are short an explanation somewhere. Speakers <u>do</u> have explicit linguistic knowledge, this knowledge is <u>not</u> justified in terms of any tacit premisses from which it was inferred

(according to Kitcher), and it is not directly evident. How then can we explain it?

Kitcher's account. Kitcher argues that we cannot answer this question satisfactorily as long as we hold on to the justification condition on knowledge explanations. This condition is a direct outgrowth of the theory that identifies knowledge as justified-true-belief (plus perhaps something more). If this characterization is correct, then it is natural to demand that an explanation of a person's knowledge state will have to show that the believer is justified in his belief. If we jettison this condition, then Kitcher believes there is no problem with linguistic intuitions that the ability approach cannot solve. Unfortunately, if we do jettison it then we will not be able to distinguish an explanation of how it is that someone knows something from an explanation of how they came to believe it. So we can't just junk the justification condition without introducing something else to replace it. It is at this point that Kitcher suggests we return to the ability view rejected earlier, and look to the causal theory of knowledge to provide an adequate replacement for the justification condition.

What Kitcher gets out of the (the attempts to develop a) causal theory of knowledge is the concept of knowledge without justification. The by now familiar idea is that besides the acclaimed Gettier cases where we have justified true belief but no knowledge, there are also cases where we have knowledge without justification. Goldman has argued that (e.g.) a chicken-sexer's knowing whether the animal in front of him is male or female is an example of such knowledge without justification. What is supposed to distinguish these cases from instances where we have no more than true belief is that, as Kitcher puts it, "the technique or causal mechanism which leads to the formation of knowledge (is) reliable".[27]

This condition lends itself very easily to our problem--viz. how do we explain the existence of linguistic intuitions? If we construe knowledge of language along the lines of the ability approach, then we can understand the speaker's explicit beliefs as resulting from a set of dispositions. These dispositions, in turn, form a causal network that reliably generates true beliefs. If we embrace the causal theory then we can explain the status of grammatical intuitions as knowledge without the postulation of any underlying tacit knowledge.

> A full explanation of the speaker's knowledge would thus proceed along the following lines. We should first describe the speaker's ability as a set of psychological dispositions, that

> is, dispositions to pass from one type of psychological state to another type of psychological state. We should then show that the psychological dispositions are regularly involved in the production of true beliefs, by describing the ways in which true linguistic beliefs are formed. (Typically, a stimulus will put the speaker into one state; in virtue of one of his psychological dispositions he will pass into another state; from this latter state, in virtue of another psychological disposition, he will pass into yet another state; and so it goes on, until the causal sequence culminates in the formation of a true belief.) Our final explanatory task would be to show that appropriate psychological dispositions were involved in the production of a true belief on the occasion in question.[28]

To sum up, Kitcher's attack has two parts. First, he argues against the viability of the Graves-Katz explanation of explicit linguistic knowledge in terms of tacit knowledge of general rules. Second, he suggests an alternative explanatory framework that purportedly does account for this explicit knowledge. In what remains of this chapter I will argue that Kitcher's theory is unsatisfactory on both counts. I begin with the alternative explanatory model.

Problems for Kitcher. (1) We should notice straight away that Kitcher's dispositionalist theory is closely related to Stitch's perceptual model. It is no accident that the causal theory Kitcher embraces has been spelled

out in greatest detail in the case of perceptual knowledge. A significant consequence of this relation is that Kitcher's model is subject to the same problems we found in Stitch's account. To repeat our earlier point, when we judge the grammatical status of a sentence we do not discover some new fact about the world--we are reporting something that we, in a sense, 'knew all along'. In this regard at least, the grammatical and perceptual cases are disanalogous, and it seems to me that the dispositionalist approach does not take account of this fact. Perhaps the force of the disanalogy can be captured by saying that in the language case exposure to the relevant expression is not part of the cause of our knowing that the expression has/doesn't have the property in question. We already have that knowledge; being exposed to the relevant expression is simply the occasion of our doing the necessary processing to become aware of it. In perception-type cases, a causal model seems more appropriate. To see this, consider the case of a baseball umpire who is asked whether a batter would be 'out' in such-and-such a situation. It would be misleading to characterize his hearing our description of the situation as (part of) the cause of his knowing whether the batter would in fact be 'out'. It is more natural to treat his answer as the result of an inference from a set of known rules. Like the speaker, the umpire

is not discovering anything new when he makes his judgement.

(2) If the dispositionalist approach is to pose a serious challenge to the Graves-Katz model, then it must be shown that it is as simple or simpler than the inferential model. Kitcher explicitly claims that it is. I want to explain briefly why there is reason to believe that it is not. The source of the problem is that the dispositionalist model must provide a separate 'mini-theory' for every independently specifiable ability or task that involves our knowledge of language. Consider explicit knowledge of grammatical properties for a moment. It should be obvious that there are any number of language related abilities that involve this explicit knowledge. To take a simple example, besides our being able to answer direct questions about the properties of expressions, Kitcher himself has pointed out that we use such knowledge when we assess the soundness of deductive arguments (recall the Jimmy Carter example). A dispositionalist account would have to explain each of these abilities--to supply direct intuitions, to spot (in)valid inferences, to correct errors, etc.--in terms of a separate network of dispositions that mediate between the (obviously different) stimuli in each case, and the (again different) responses that the subject could be

expected to make. In fact, the proliferation would seem to be unbounded given the range of phenomena to be explained--for instance, the tendency to displace clicks to syntactic phrase boundaries. For the inferential model, all these capacities tap one knowledge source. Of course, the dispositionalist could argue that what all these behavioral capacities have in common is not some body of tacit knowledge but rather an interlocking subset of dispositions which is activated and plays a part in all the causal chains that subtend these capacities. But this begins to sound like a tacit knowledge explanation. How can we characterize this common subset except to say it involves the speaker's knowledge of general rules?

This should recall our earlier discussion of dispositional analyses. A listing of the input-output (in this case, stimulus-response) relationships of the subject only helps to organize the problem; it does not constitute a solution. We still have to ask what features of the subject account for these correlations--why he has these dispositions. Now it could be that Kitcher believes, following Quine, that these dispositional analyses will one day be replaced by neurophysiological explanations. This is obviously a complex issue that we cannot do justice to here. Let me simply say that I know of no compelling reason to believe that the ultimate theory that will account for the fact

that speakers have linguistic intuitions will be statable in terms of neurophysiological primitives. But even if there is such a reason, it is not clear what bearing it should have on ascriptions of tacit knowledge. This brings us to our next point.

(3) In one sense the whole attempt to provide dispositionalist approaches of mental states, right or wrong (and I believe it is wrong), is irrelevant to the problem of ascriptions of tacit knowledge. Let me explain why.

Consider the everyday, uncontroversial notion of belief--e.g. my belief that the bus will be late. The classical dispositionalist account has it that to believe that the bus will be late is to be disposed to act in certain specifiable ways given the appropriate stimuli. Let's say we ignore the by now familiar shortcomings of this account, and we ask a philosopher who favors this sort of analysis whether people have beliefs, or whether they only have dispositions. This sort of question clearly misses the point; who in their right mind would deny that people have beliefs? The point of the analysis is to free the concept of belief from its mentalist associations; what we ascribe when we ascribe a belief to someone is a set of behavioral dispositions. If there

were never any beliefs to start with there would be very little for dispositionalists to talk about. What the theory tries to do is analyze the concept of belief--not eliminate it.[29]

In light of this, it is somewhat surprising that Kitcher tries to use the dispositionalist approach to eliminate the notion of tacit knowledge. He seems to say that since we can explain explicit linguistic knowledge in terms of dispositions to form the relevant beliefs, we no longer need to talk about any underlying tacit knowledge. But why should this be so? It is no doubt true that we are disposed to form true beliefs when queried about the grammatical status of expressions in our language. But why should this obviate the need for an explanation in terms of underlying knowledge of general principles?

Consider the following example. Sherlock Holmes and Dr. Watson are standing over a corpse at the scene of the crime and Holmes suddenly announces that the murderer has studied surgery. Holmes rightly claims to know this for a fact. He has deduced it from the fact that there is a mark of a perfectly placed surgical needle at the base of the skull. Because of his years of study and experience, Holmes arrives on the scene with a store of world-knowledge which includes: that a perfectly placed

needle can kill a man, that the position of this puncture mark coincides with one such place, that only one who has studied surgery could know how to administer the injection, and so on.

The 'received view' of this incident (what Conan Doyle would want us to accept) is as follows. Holmes knows three sorts of things here:

(1) his store of world knowledge,

(2) that the puncture mark on the victim's body is at the base of the skull, the the person is dead, etc.,

(3) that the perpetrator of the crime has studied surgery.

Put in our terms, Holmes knows (1) explicitly. These are generally long-standing beliefs that can be brought to the fore as is necessary. He also knows (2) explicitly, but this is knowledge of recent vintage that has its source in present perceptions; these are things Holmes just found out. (3) is also explicit knowledge of recent vintage, but it is not perceptual knowledge. The process by which Holmes comes to know (3) is naturally construed as inferential-- "Deduction (sic), my dear Watson".

Now imagine that a dispositionalist offers a variant epistemological report about this incident: Holmes came upon the scene of the crime and was exposed to a set of stimuli S1...Sn (the puncture marks, etc.) and this caused him to be in psychological state P1. In virtue of

one of his psychological dispositions he passed into psychological state P2. From this state, in virtue of another disposition, he passed into state P3. Eventually, the the causal sequence terminated in the formation of the true belief (or: the utterance to the effect that) the perpetrator studied surgery.

This parody of Kitcher's account certainly lacks the snappiness of Holmes' own explanation ("Deduction!). But snappiness aside, are the two accounts incompatible? Is the dispositionalist's story meant to replace the 'received view'? Doesn't everyone grant that Holmes knows (1)-(3) and that he inferred (3) from premisses that include (1) and (2)? It would seem that the answer here is 'yes'. So the fact that we can formulate an account (which I have argued is inadequate) in terms of dispositions doesn't seem to clash with the ascription of knowledge states and inferential processes.

Part of the problem is that Kitcher often talks as if the tacit knowledge model tries to explain linguistic intuitions in terms of non-causal processes. This is an error. The claim that (a) every instance of a linguistic intuition is the end product of a causal process that is mediated by a reliable mechanism, does not in any way challenge the complementary claim that (b) the causal sequences that mediate <u>this</u> relationship should be cha-

racterized as inferences from general principles. To avoid misunderstanding, let me stress that I am not arguing that all causal processes that could mediate this relationship must be characterized as 'making inferences from general principles'. There is no necessity here; it is empirically possible that the reliable mechanism in this case is not inferential. But the causal theorist does not claim that knowledge is never the result of inferential processes--he only wants to say that it need not be. What I claim is that in this case, the inferential model is the most plausible hypothesis.

(4) We might note in passing that the version of the causal theory of knowledge that Kitcher adopts seems as vulnerable to Gettier-type counterexamples as is the justified-true-belief theory. Following Goldman, Kitcher counts true belief as knowledge if "the technique or causal mechanism which leads to the formation of knowledge (is) reliable". But now consider the following example. Let's say we have a reliable mechanism--one which almost always leads to true beliefs--and that in a certain case it leads to the belief that p.[30] Now it could be the case that p is true, but that in this specific instance the reliable mechanism did not do its job. Still, some fluke interference (a shower of gamma rays?) interacted with the reliable mechanism and contributed to

the formation of the true belief. If Kitcher's formulation is right, then this belief should count as knowledge; but our intuitions tell us otherwise.[31]

(5) Finally, leaving aside the general issues about the status of dispositional analyses and the causal theory of knowledge, it seems to me that Kitcher's original argument--i.e. the nativist's 'dilemma'--makes an unwarranted assumption. The 'dilemma', we will remember, turns on the impossibility of providing a justification for innate knowledge of a universal grammar. Kitcher argues that since such knowledge is not to be justified on the basis of reasoning from any prior knowledge, it must be 'directly' justified. This much seems correct. But he then goes on to reject the idea that innate knowledge can be directly justified, since it is not at all like the paradigm of directly justified knowledge--viz. perceptual knowledge. If we disregard the paradigm, then why not say that intuitions are directly justified?[32]

But here I would argue that Kitcher again makes a mistake closely related to the one we found in Stitch's argument. He in effect <u>assumes</u> that the only beliefs that can be directly justified are those that are directly evident. What reason is there to think that this is so? It has been argued that we can legitimately take

the inductive principle to be directly justified while at the same time deny that any particular inductive inference is directly justified. Similarly, one might hold that the general innate principles of universal grammar are directly justified but deny that such justification extends to any specific linguistic intuitions. I treat these questions briefly here, because they are outside the limited scope of the problem of tacit knowledge. They are more pertinent to an analysis of the relation between rationalism, innateness, and knowledge, and to the problem of the foundations of knowledge. Still, there seems to me to be a good case for the claim that certain general principles of thought are instances of knowledge even though they are neither indirectly justified nor directly evident.[33, 34]

In summary, I have argued that Kitcher's 'dilemma' is based on an undefended conception of the possibilities of direct justification. As far as his positive theory is concerned, I have tried to show that the strategy of joining the dispositionalist analysis of belief to the causal theory of knowledge has very little to recommend it.

NOTES CHAPTER IV

1. "Tacit Knowledge", JP, LXX, 11, June 7 1973, pp. 318-330.

2. The point has been noticed, but its significance has been overlooked. Harman, for instance, grants that speakers typically have such knowledge (which he calls 'tacit') but he points this out to contend that this is the only sort of propositional knowledge speakers have. See his "Psychological Aspects of a Theory of Syntax" in Stitch, Innate Ideas, pp. 172-173, and Thomas Nagel's "Linguistics and Epistemology" in Hook, Language and Philosophy, p.174.

3. Again, the specific examples are not meant to signal a commitment to any specific theory about the sorts of underlying rules speakers use.

4. "Tacit Knowledge", p. 329.

5. Ibid., pp. 325-326.

6. See Chapter I.

7. "Tacit Knowledge", p. 329.

8. See the discussion of Nagel's position in Chapter II for a similar confusion.

9. This should not be taken to imply that there is a clear community-wide agreement about what is and what is not to count as grammatical.

10. It is admittedly difficult to imagine this turning out to be the case with something like (1).

11. As Graves-Katz point out, there is some confusion here between 'grammatical in English' and 'grammatical in my idiolect'.

12. Stitch, "What Every Speaker Knows", p. 478.

13. Hilary Putnam, "The 'Innateness Hypothesis' and Explanatory Models in Linguistics", in The Philosophy of Language, ed. John Searle, (London: Oxford University Press, 1971), pp. 131-132.

14. "Tacit Knowledge", p. 319n.

15. "What Every Speaker Knows", pp. 492-493.

16. Philip Kitcher "The Nativist's Dilemma" Philosophical Quaterly, v. 28, no. 110, January 1978, p. 6. For a discussion of related issues see the "Introduction" to The Logic of Grammar, ed. Donald Davidson and Gilbert Harman, (Encino, California: Dickenson Publishing Co., Inc., 1975.

17. See Chapter I.

18. "What Every Speaker Knows", p.493.

19. Ibid., pp. 493-494.

20. Ibid., p. 495.

21. Ibid.

22. Not all nativists agree that such a move is required, and one can, alternatively, defend the notion of tacit knowledge and steer clear of the nativism issues. In this work I take a position close to this. A prime example of a nativist who takes the first line is Chomsky himself, who in recent writings has been willing to bypass the question of tacit knowledge and try to make the case for nativism in terms of a new notion--'cognizance'. See in this regard Chomsky's Reflections on Language.

23. "Tacit Knowledge", p. 318.

24. "The Nativist's Dilemma", pp. 6-7.

25. I assume of course that p is true.

26. Ibid., p.6.

27. Ibid., p. 7n.

28. Ibid., pp. 8-9.

29. I don't want to give the impression that this is a well-understood distinction--it isn't. Still, only the most radical anti-mentalists would claim that there are really no beliefs. Secondly, there is a difference between accepting the fact that people have beliefs and reifying beliefs. On this see Arthur Collins "Could Our Beliefs Be Representations in Our Brains?" JP, LXXVI, no. 5, May 1979, pp. 225-243.

30. This is not intended to be a precise explication of 'reliability'.

31. David Shatz suggested this line of attack to me.

32. This line of thought is closely related to Harman's argument taken up in the first part of Chapter III.

33. For an interesting foray into these questions see Nagel's "Linguistics and Epistemology", section II.

34. This point was suggested to me by Jerry Katz's personal response (correspondence) to Kitcher's argument.